DOCUMENT RESUME

ED 207 587 IR 009 705

AUTHOR Carr, Brian
TITLE Wusor II: A Computer Aided Instruction Program with Student Modelling Capabilities. AI Memo 417.
INSTITUTION Massachusetts Inst. of Tech., Cambridge. Artificial Intelligence Lab.
SPONS AGENCY National Science Foundation, Washington, D.C.
REPORT NO LOGO-45
PUB DATE May 77
NOTE 133p.; Master's Thesis, Massachusetts Institute of Technology.

EDRS PRICE MF01/PC06 Plus Postage.
DESCRIPTORS Artificial Intelligence; Computer Assisted Instruction; *Computer Programs; Decision Making Skills; *Games; Logical Thinking; *Models; *Problem Solving; *Tutoring
IDENTIFIERS *Computer Games; *Intelligent CAI Systems; Tutorial Mode; Wumpus; Wusor

ABSTRACT

Wusor II is the second intelligent computer aided instruction (ICAI) program that has been developed to monitor the progress of, and offer suggestions to, students playing Wumpus, a computer game designed to teach logical thinking and problem solving. From the earlier efforts with Wusor I, it was possible to produce a rule-based expert which possessed a relatively complete mastery of the game. Wusor II endeavors to teach the knowledge embodied in the rules used by the expert. The student model represents Wusor's estimation of the student's knowledge of these rules, and this estimation is based primarily on analyses of the player's moves. The student model allows Wusor to personalize its explanations to the student according to the student's current knowledge of the game. The result is a system which, according to preliminary results, is highly effective at tutoring students of varied abilities. Thirty-three references are listed. (Author/LLS)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

IR

ED207587

Artificial Intelligence Laboratory
Massachusetts Institute of Technology

AI Memo 417 Logo Memo 45

Wusor II:
A Computer Aided Instruction Program With Student Modelling Capabilities

U S DEPARTMENT OF HEALTH
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

Brian Carr

May 1977

Abstract

Wusor II is the second program that has been developed to tutor students in the game of Wumpus. From the earlier efforts with Wusor I it was possible to produce a rule-based expert which possessed a relatively complete mastery of the game. Wusor II endeavours to teach the knowledge embodied in the rules used by the Expert. The Student Model represents Wusor's estimation of the student's knowledge of said rules, and this estimation is based primarily on analyses of the player's moves. The Student Model allows Wusor to personalize its explanations to the student according to the student's current knowledge of the game. The result is a system which, according to preliminary results, is highly effective at tutoring students of varied abilities.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. It was supported in part by the Division for Study and Research in Education, Massachusetts Institute of Technology and has been submitted to the Department of Mathematics in partial fulfillment of the requirements for a Master of Science Degree under the supervision of Professor Ira Goldstein of the Department of Electrical Engineering/ Computer Science.

IR009705

Table of Contents

Table of Contents 2

Chapter 1 .... Introduction 3

Chapter 2 .... The Original Wumpus Advisor 7

Chapter 3 .... The Expert 13

Chapter 4 .... The Improved Wumpus Advisor 29

Chapter 5 .... The Psychologist and Student Model 38

Chapter 6 .... Tutoring Strategies 53

Chapter 7 .... The English Generation Routines 60

Chapter 8 .... Testing The Advisor 65

Chapter 9 .... Conclusion 80

Appendix A .... Detailed Modular Diagram 84

Appendix B .... Details of Move Comparisons 88

Appendix C .... Ideal Situations, Descriptions 93

Appendix D .... Ideal Situations, Examples 95

Appendix E .... Ideal Wumpus Scenario 99

Appendix F .... Results From Synthetic Student Tests ... 107

Appendix G .... Sample Synthetic Student Scenario 111

Bibliography 129

3

The Introduction

Chapter 1

One of the purposes of our educational system is to teach better ways to think about problems in general. This is a key justification for the study of higher mathematics and the sciences by the general student. Similarly, games can be an ideal teaching instrument. Most games represent some facet of life but in a much simpler domain, thereby allowing the student to develop his problem solving abilities in an environment which is more conducive to learning. Such games can be fun while teaching valuable skills. However, students can reach a plateau in their game playing and cease to try new strategies. When this occurs, the learning process halts. A solution to this problem is to encourage the student to improve by occasionally offering suggestions regarding improved game strategies. Unfortunately, the cost of providing human teachers to watch the game (and offer suggestions) is prohibitive. However, it is possible to use a computer to monitor the progress of the game and to offer suggestions when warranted. With this goal in mind, a program was written to serve as just such an advisor for a computer game referred to as "Wumpus hunting" or "Wumpus" (1).

The relationship of the Wumpus Advisor to the player is illustrated in Figure 1.1. The Advisor analyzes the interaction between the game and the player, giving advice when appropriate. When the Wumpus Advisor chooses to give advice, it can intercept the player's messages to the game and allow the player to modify his move if he so desires. Also, from the player's point of view, the Wumpus Advisor receives exactly the same information as the player. The Wumpus Advisor can alter the game (shown in dotted lines) to a certain degree. It does so to create situations which are more conducive to learning.

---

(1) The game of Wumpus was originally described by Yob in Creative Computing, and this particular variation was implemented in LISP by Greg Clemenson.

Figure 1.1
Conceptual Model of a Tutor.

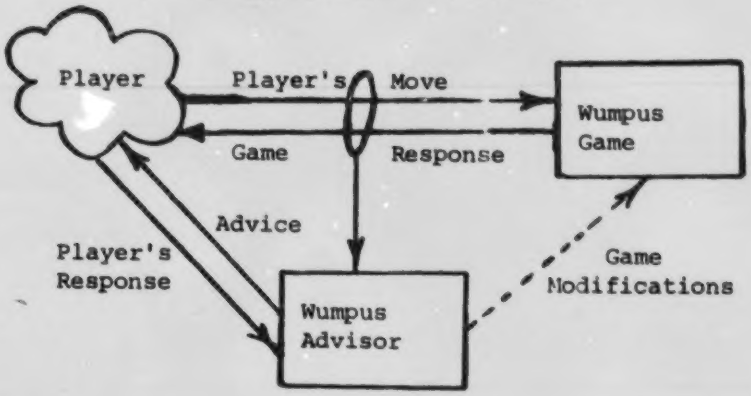

The game "Wumpus" is an example of a game which encourages deductive processes and develops a sense of probabilities. The player must seek out and kill a monster, the Wumpus. The player moves in the Wumpus's warren, a network of interconnecting caves containing the Wumpus and other dangers, namely bats and pits. At the start of the game the player is told the number of caves in the warren and the number of bats and pits. (ii) Before every move the player is told which cave he is in and the neighboring caves that he can move to. (iii) If any of the neighboring caves contain a bat he will be informed that he hears

---

(ii) In a normal game, there are twenty caves, three pits, and three caves with bats.

(iii) Each cave is identified by number.

squeaking. Likewise, if a neighboring cave contains a pit, the player will be informed of a draft (as pits are bottomless chasms), but in neither case is the player told which of the caves are dangerous. Whenever the player is within two caves of the Wumpus he will smell its horrible stench. If the player enters a cave with a pit he loses the game, whereas if he enters a cave containing a bat, he will be carried to a random cave which may contain another bat, a pit, or the Wumpus. The Wumpus eats unwary players who stumble into his lair. The player tries to visit enough caves (avoiding bats and pits) to locate the Wumpus without actually entering its cave. Once the player has found the Wumpus, he can shoot an arrow into the Wumpus' lair from a neighboring cave, killing the beast. If the player shoots an arrow into a cave and the Wumpus is not there, his arrow will ricochet through the warren at random for roughly four caves and may kill either the player or the Wumpus. In Figure 1.2 part of a possible warren is shown.

Figure 1.2

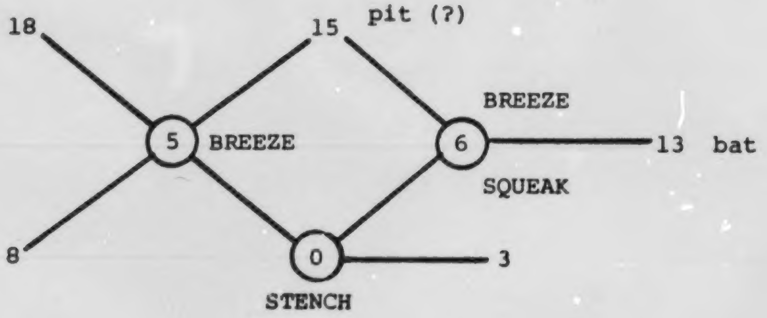

Circled caves have been visited.
Upper Case: Given Facts
Lower Case: Deductions

Playing the game can involve simple deductions and risk

C

minimizations as well as more complex strategies and considerations. At the simplest level, the player can deduce that certain caves are *absolutely safe* by the absence of warnings. These caves should be explored before any others. At a higher level the player can perceive that certain caves are *probably safe*, i.e. less likely to contain dangers than other caves. These perceptions are based on the patterns of the warnings and require application of probabilistic heuristics which are commonly used by knowledgeable persons. An advanced player can usually deduce the exact location of the Wumpus through a quite thorough application of logic, but most players develop a general idea of its location without completing all the required deductions (though they are wrong often enough to encourage them to improve their deductive powers). There are also unusual situations which require very advanced considerations of the risks involved in order to select the best move. The Wumpus is a game that can be enjoyed by the beginner as well as the advanced player.

In the development of the Advisor, two different generations of programs have been written. They are referred to as Wusor I (Wumpus Advisor I) and Wusor II. For the sake of perspective, the next chapter will discuss the development of Wusor I, and then the remainder of the paper will discuss Wusor II, the most recent Wumpus Advisor.

The Original Wumpus Advisor, Wusor I

Chapter 2

Wusor I can be viewed as a continuation of the work Burton and Brown (1975) with the game of West. West is an algebraic game for which they developed an Advisor program. However, West was a relatively simple domain which was not conducive to more advanced Advisors. As a result, it was decided that Wumpus was a domain that was sufficiently difficult to allow the development of more advanced Advisors, though it was still simple enough to allow the development of a functional program. With this goal in mind, work was begun on Wusor I.

Wusor I was written as a course project for nine graduate students at M.I.T. with its development being supervised by Jim Stansfield and Ira Goldstein. As the Wumpus Advisor was required to discern the quality of the different moves as well as be able to explain itself when appropriate, it quickly became a very complex program. To keep everything manageable, the program was written in a modular fashion. This aided in the development of the Advisor as it allowed the modules to be written, debugged, and modified relatively independently. After its initial development, Wusor I was extensively modified by the author. [i]

Wusor I had three main modules and several sub-modules. The three main modules were the Executive, the Bat/Pit Locator, and the Wumpus Locator. [ii] Associated with the three main modules were various sub-

---

[i] Wusor I is fully described in Stansfield, Carr, and Goldstein (1976).

[ii] The Bat/Pit Locator was written with the help of Will Clinger, John Avgoustis, and Fred Knowle. The Wumpus Locator was originally written by

Figure 2.1
Dotted lines represent significant flows of information.
Solid lines illustrate the hierarchy of control.

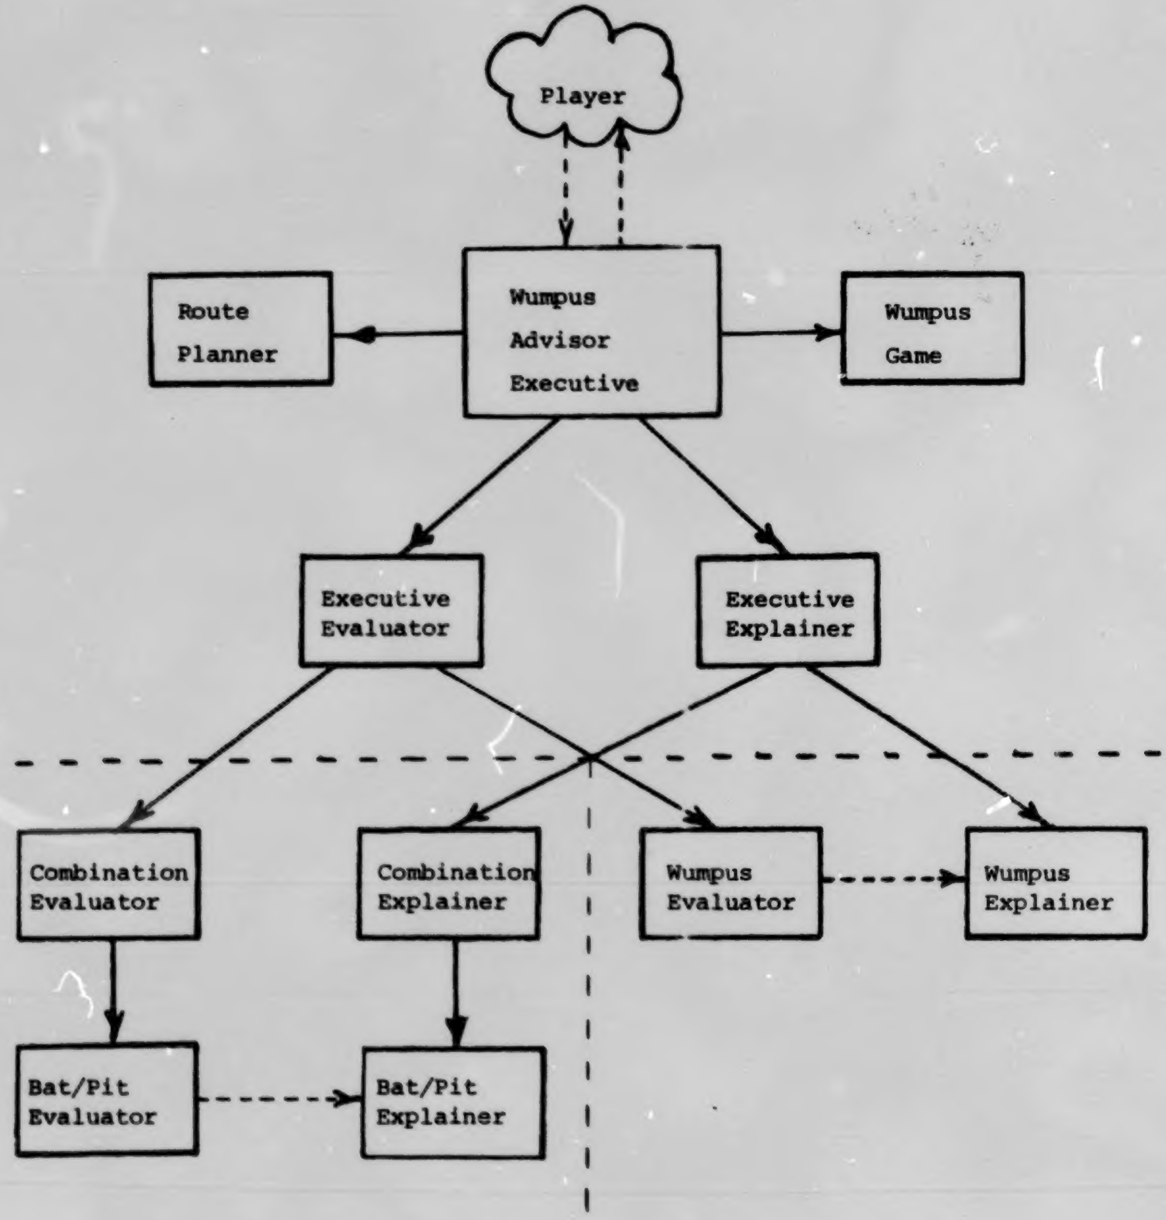
Player
Route Planner
Wumpus Advisor Executive
Wumpus Game
Executive Evaluator
Executive Explainer
Combination Evaluator
Combination Explainer
Wumpus Evaluator
Wumpus Explainer
Bat/Pit Evaluator
Bat/Pit Explainer

modules which are illustrated in Figure 2.1. Each of the Locator modules had two sections, the Evaluator section which computed the merits of the various caves, and the Explainer section which explained the results of the Evaluator. Also, the Bat/Pit Locator had two levels, the level which considered bats and pits separately, and the Combination level which combined these results. Finally, the Executive Evaluator combined the results from the Bat/Pit Evaluator and the Wumpus Evaluator to come up with an overall utility for each cave. Then the Executive had the option of calling upon the Explainer routines to generate explanations for the student. This structure is rather contorted and unnecessarily hard to understand, but this is primarily due to an incomplete understanding of the domain at the start of the project. However, from the work on Wusor I it was possible to gain many insights which finally led to the development of Wusor II.

The Bat/Pit Locator computed an estimate of the probability that any cave contained a bat or a pit, but the Wumpus Locator was not advanced enough to arrive at estimation of the risks from the Wumpus; it only determined which caves could not contain the Wumpus and which caves were likely to return information about the Wumpus' location. (iii) The Executive Evaluator could not directly compare the risks involved from Bats/Pits and the Wumpus as it did not have any probability estimates for the Wumpus. Instead it broke down the different moves into the eight categories shown in Figure 2.2 and evaluated the player's move according

---

Neil Rowe, Beth Levin, and Robin Gross. The Route Planner was written by Ginny Grammar.

(iii) The Wumpus Locator was sufficiently advanced to locate the Wumpus, but it could not deal with incomplete information.

to these categories. The different Evaluator routines taken together comprised the Wumpus "Expert" (iv) which possessed a knowledge of the game.

Figure 2.2
Wusor I uses eight classifications for unvisited caves.

Type 1. *A cave that is absolutely safe and that gives information about the Wumpus.*

Type 2. *A cave that is absolutely safe.*

Type 3. *A cave that has no bats or pits, but that could contain the Wumpus. This cave inherently gives information about the Wumpus.*

Type 4. *A cave that might contain bats or a pit, but not the Wumpus and which gives information about the Wumpus.*

Type 5. *A cave that might contain bats or a pit, but not the Wumpus.*

Type 6. *A cave which could contain bats, pits or the Wumpus.*

Type 7. *A cave that contains a bat. This classification is a very special class that has varying value.*

Type 8. *A cave that is certain death. (Either a pit or the Wumpus).*

Throughout the development of the Expert, heuristics were always used in preference to lengthy calculations. This was partly because the exact calculations were so lengthy that it was unrealistic to compute them, but also because such heuristics were, in fact, used by humans players. (v) This correspondence between human and computer reasoning greatly simplified explanations. An example of this was the manner in which the Bat/Pit Locator determined the likelihood that a cave contained

---

(iv) The concept of designing the tutor around an expert was developed by Brown (1973) in his Sophie program.

(v) In particular, the heuristics were those used by the designers of Wusor I. Further experimentation is required to determine how common these heuristics are.

a bat or a pit. It might have been possible to compute the exact probability of a cave containing a bat or a pit, but this calculation would have been time consuming and complex. Worse yet, any explanation of an exact probability would have had to discuss the relevant permutations and combinations, which would almost certainly have been above the level of the student.

Wusor I had only a very limited model of the user's knowledge. It did not keep track of different players, but, instead, asked the user at the start of each session for an estimation of his ability. For the rest of the session the Expert would make all of its computations at a level of complexity hoped to be slightly above that of the student. This had the disadvantage that, if the player entered an unrealistic evaluation of his ability, he would either be overwhelmed by material well above his level, or he would find the Advisor's advice over simplified. There was also the problem that Wusor I did not have fine gradations in the level of its advice (it had four levels) due to the expense of programming the Expert to function at more levels.

Wusor I was tested on a very limited basis and was found to be reasonably effective at teaching better Wumpus strategies. However, it was mostly used by M.I.T. related personnel (including children and graduate students) who, we assume, already possessed a logical bent. In these cases it was effective, but it was not attempting to teach a better way to think about problems (as the players already possessed the desired thinking habits) but merely better ways to think about a particular problem. In those very few cases where it was advising students who did not already have the desired thinking habits, the handicaps of Wusor I were readily apparent. It was not able to carefully guide the player through the more basic steps (its simplest level was not basic enough) and tended to confuse and frustrate the student. Because of these limitations, work was started on a more advanced Wumpus Advisor with a more detailed student model and the ability to tune its explanations to the particular student.

The Expert

Chapter 3

Overview of the Expert

From the work which was done on the expert of Wusor I, it was possible to develop a much more powerful and general Expert. The deductions required in locating the dangerous caves and the heuristics necessary to arrive at an estimation of the danger in cases of uncertainty have been simplified to the point where there is a Locator Algorithm which requires only very general information (such as the number of caves, the number of bats, how far a bat warning propagates, etc.) to determine the relative dangers of each cave. The Locator Algorithm has three basic steps as shown in Figure 3.1. It is executed once for each danger being considered, and the results are then combined probabilistically. The first step is the acceptance and propagation of any and all information supplied by the Wumpus game.

Figure 3.1
The Three Steps of the Locator Algorithm

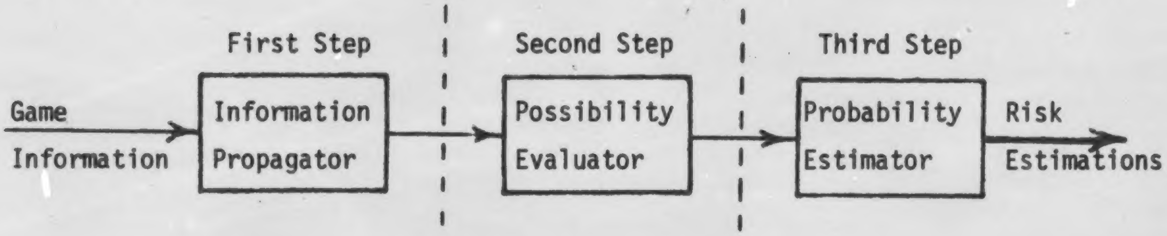
First Step
Second Step
Third Step
Game Information
Information Propagator
Possibility Evaluator
Probability Estimator
Risk Estimations

The Information Propagator.

For each of the three dangers the Information Propagator marks each cave with three pieces of information, the maximum distance to the closest danger, the minimum distance to the closest danger, and, when the first two numbers eliminate all other possibilities, a marker for the exact distance to the closest danger. There are eight theorems which allow the Propagator to update this information about a given cave. The first two are that any time a cave is visited it can be marked as either *more than zero away* or exactly *zero away.* (i) Also, in the special case of the Wumpus, if the player shoots an arrow into a cave and does not kill the Wumpus, then that cave is *more than zero away.* (ii) Next, if a cave is visited and has no warning, then that cave can be marked as *more than N away*, where N is the distance that a warning for the given danger travels. Likewise, if a warning is sensed, then the cave must be *less than (N+1) away.* The last theorems deal with the propagation of information through the warren. If a given cave is *more than N away*, then all of its neighbors must be *more than (N-1) away.* Conversely, if all of a cave's neighbors are *more than (N-1) away*, then said cave can be classified as *more than N away.* The above theorems correspond to rules zero through seven shown in Figure 3.2. Using additional theorems, more information could be propagated, but these eight rules are sufficient to mark each cave with all the information relevant to playing the game. Many caves can be classified as being safe just through the use of the *more than N away* tag. There are examples of the deductive process in Figures 3.3 and 3.4.

(i) Except for bats, the classification of *zero away* attained in this fashion is rather pointless as it also signifies the end of the game.

(ii) Note that in some implementations of Wumpus, the Wumpus moves when an arrow is shot (if it does not kill him). Wusor II does not permit this primarily because it would require degrading the data base. This could create situations which were too complicated to readily be explained to children.

Figure 3.2
The Logical Rules of the Expert

L0 *A cave can be marked as "zero away" if it was visited and found to contain a danger.*

L1 *A cave can be marked as "more than zero away" if it was safely visited.*

L2 *If the player shoots an arrow into a cave and does not kill the Wumpus, then that cave can be marked as "more than zero away" (Wumpus).*

L3 *If a cave is visited and there is not a warning, then that cave is "more than N away" where N is the distance that the warning propagates.*

L4 *If a cave is visited and there is a warning, then that cave is "less than (N+1) away".*

L5 *If a cave is marked "more than N away" then all of its neighbors must be "more than (N-1) away".*

L6 *If all of a caves neighbors are "more than (N-1) away", then it must be "more than N away".*

L7 *If a cave is "more than (N-1) away" and "less than (N+1) away", then it is "N away".*

L8 *When the algorithm is creating cave-sets and it encounters a cave which would be N caves away but which is also "more than N away", then that cave can not have any contributions to the cave-set.*

L9 *If the player encountered a danger in a cave, then the cave does not contain a danger of higher priority, i.e. the Wumpus eats the player before he can fall into a pit, and he will fall into a pit before he is picked up by bats.*

L10 *Certain caves can be marked as "more than zero away" based on consideration of the different complete cave-sets and the number of dangers.*

In Figures 3.3 and 3.4 there is a standard notation for diagrams of the warren. The nodes of the network are numbered and represent caves. Circled numbers represent caves which have been visited by the player. To the top right of each visited cave is a marker for whether or not any warnings were sensed. ("W" indicates that a warning was sensed, and a "NW" means that a warning was not sensed) To the bottom left of each cave is any classification that the cave may have. (i.e. *more than N*, *less than M*, and *J away*) The reader is cautioned that, while the the diagrams of the warren in the examples are very regular (planar), normally warrens are much more random in their arrangement. (iii)

In Figure 3.3, the danger being considered is pits. Hence, warnings propagate for a distance of one cave. In Figure 3.3A, the player has visited two caves, neither of which had any warnings. Therefore, the visited caves are *more than one away* and all the neighbors are *more than zero away* (i.e. safe). Upon visiting cave 3 (Figure 3.3B), the player receives a warning, indicating that cave 3 is *less than two away*. Since it is also *more than zero away*, it must be *one away* (i.e. cave 3 is next to a pit).

In Figure 3.4 the danger is the Wumpus whose warning propagates for a distance of two caves. At the start (Figure 3.4A), the player has visited various caves and has discerned that caves 7 and 10 must contain pits. He has also visited two caves of relevance, caves 5 and 6. As no warnings of the Wumpus were sensed, caves 5 and 6 were marked as *more than two away*. All the neighboring caves were then marked as *more than*

(iii) An interesting experiment is to determine whether the restriction to regular warrens increases the speed at which people learn the game.

Figure 3.3

The Danger is Pits.
W indicates a warning.
NW indicates the absence of a warning.
A warning propagates for one cave.

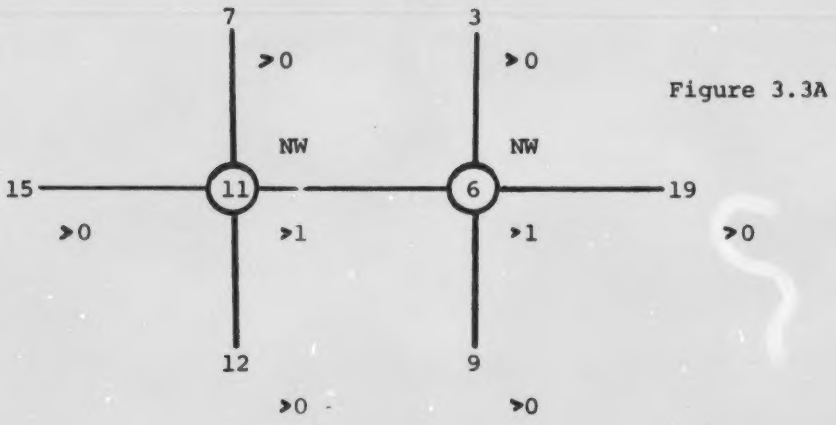
7
3
>0
>0
NW
NW
15
11
6
19
>0
>1
>1
>0
12
9
>0
>0

Figure 3.3A

Figure 3.3B

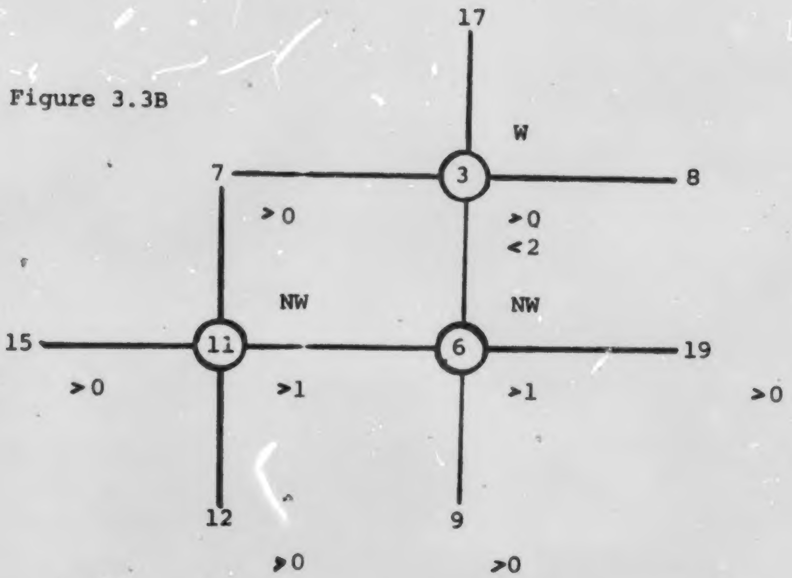
17
W
7
3
8
>0
>0
<2
NW
NW
15
11
6
19
>0
>1
>1
>0
12
9
>0
>0

one away. Next the player visited cave 3 (Figure 3.4B) and sensed a warning. This indicated that cave 3 was *less than three away*. As it was also *more than one away*, cave 3 was marked as being exactly two caves from the Wumpus. Caves 2 and 4 were marked as *more than zero away* by dint of cave 3's classification of *more than one away*. In Figures 3.4C and 3.4D the player visited caves 2 and 4 which he knew to be safe. At these caves he sensed a warning and marked the caves accordingly. Finally he visited cave 9 (Figure 3.4E) and could then conclude that the Wumpus must be in cave 1 (by a process to be described later).

The special classification of *N caves away* is made because it is the only classification used extensively in the later steps of the Locator Algorithm. In particular, if one can determine all caves that are exactly N caves away from a cave which is classified *N away* then one has what is referred to as a complete cave-set. (A partial cave-set is one in which all the caves which are known to be N caves away are listed, but the set is not necessarily complete). In the case of bats and pits, the determination of the complete cave-set is especially simple since the only classification of relevance is *one away* and the complete cave-set is just the immediate neighbors. Cave-sets are important because *at least one member of each complete cave-set must contain one of the given dangers*. Therefore one can eliminate from the cave-set those caves which can not possibly contain the given danger for other reasons (as indicated in the data base). The result after this elimination is called a reduced cave-set. (Logical rule L8 in Figure 3.2 is the formal statement of the rule to reduce cave-sets).

In Figures 3.3 and 3.4 there are several examples of cave-sets. In

Figure 3.4

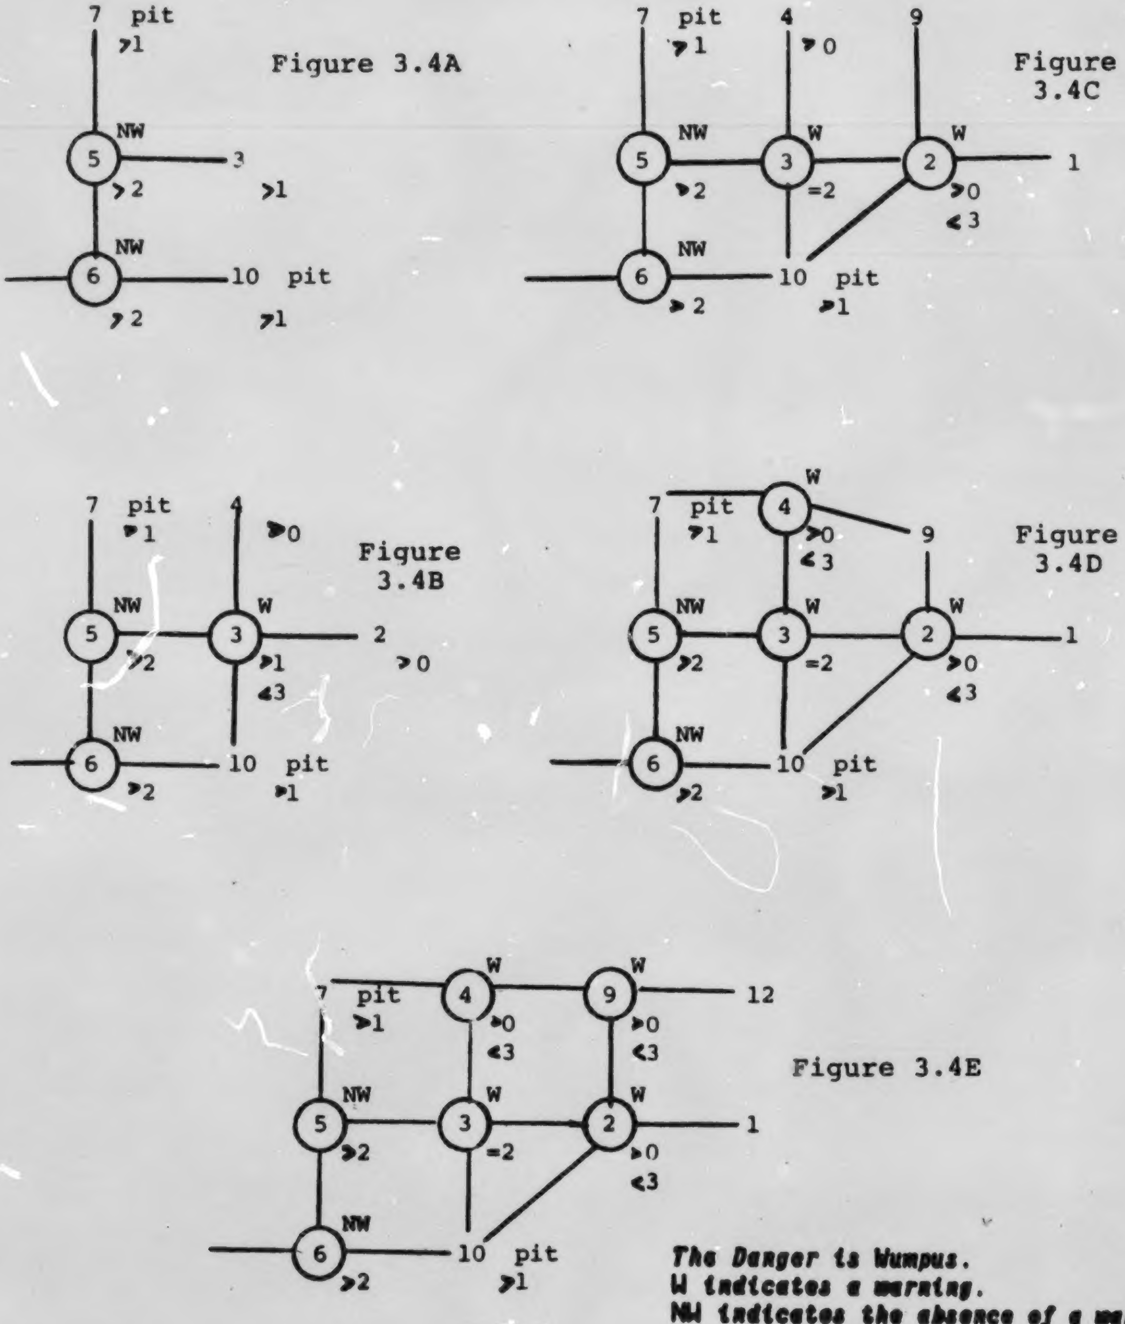
Figure 3.4A
7 pit
>1
NW
5
3
>2
>1
NW
6
10 pit
>2
>1
Figure 3.4C
7 pit
>1
4
>0
9
NW
W
W
5
3
2
1
>2
=2
>0
<3
NW
6
10 pit
>2
>1
Figure 3.4B
7 pit
>1
4
>0
NW
W
5
3
2
>2
>1
<3
>0
NW
6
10 pit
>2
>1
Figure 3.4D
W
4
7 pit
>1
>0
<3
9
NW
W
W
5
3
2
1
>2
=2
>0
<3
NW
6
10 pit
>2
>1
Figure 3.4E
W
W
4
9
12
7 pit
>1
>0
<3
>0
<3
NW
W
W
5
3
2
1
>2
=2
>0
<3
NW
6
10 pit
>2
>1

The Danger is Wumpus.
W indicates a warning.
NW indicates the absence of a warning.
A warning propagates for two caves.

Figure 3.3B the complete cave-set is caves 6, 7, 8, and 17. However, as caves 6 and 7 are both *more than zero away*, the reduced complete cave-set is just caves 8 and 17. This means that either 8 or 17 must contain a pit. In Figure 3.4C the partial cave-set is caves 6, 7, 1, 10, and 9. As caves 6, 7, and 10 are safe from the Wumpus, the reduced partial cave-set is caves 1 and 9. In Figure 3.4D there is an excellent example of the application of logical rule L8. The reduced cave-set is caves 1 and 9, but this cave-set is also a complete cave-set as the neighbors of cave 4 are known and, while the neighbors of cave 10 are not known, cave 10 is marked as *more than one away* and none of its neighbors can possibly contain the Wumpus. (Hence they can not be members of any reduced cave-set for the Wumpus). When the player visits cave 9 (Figure 3.4E) he is taking a big chance, but once he has visited it he can eliminate it from the reduced complete cave-set (it is *more than zero away*). Then the player can conclude that the Wumpus must be in cave 1 as it is the only member of the reduced complete cave-set. In all future references, it will be assumed that a cave-set is actually a reduced cave-set, since the Locator Algorithm always reduces the cave-sets.

The Possibility Evaluator

The Possibility Evaluator marks certain caves as "more than zero away" by considering the different complete cave-sets and the number of dangers. It implements logical rule L10 shown in Figure 3.2. If the number of complete cave-sets exceeds the number of dangers, then it is likely that this step of the Algorithm will deduce that certain caves can not contain a danger. For example, if there are only three caves with the given danger, and there are complete cave-sets of:

(1,2)

(3,4,5)
(6,7,8)
(8,9),

then one can deduce that caves 6, 7, and 9 must be safe. This means that cave 8 must contain one of the given dangers. The result after the application of the second step is:

(1,2)
(3,4,5)
(8).

An explanation for why cave 6 must be safe could be, "*You can deduce that cave 6 must be safe as there are only three pits. You know that there must be a pit in either cave 1 or 2, just as you know that there must be a pit in one of the caves 3, 4, and 5. Likewise, there must be a pit either cave 8 or 9. This accounts for all three pits and hence there can not be a pit in cave 6.*"

An algorithm to determine which caves are safe in this case is to choose N [(iv)] complete cave-sets that have no intersections (i.e. no members in common). Then any caves that are not a member of the union of the N cave-sets must, of necessity, be safe. The selection of the N complete cave-sets is carried out by exhaustively trying all possibilities.

Unfortunately, this algorithm does not find all possible caves that must be safe. For example, if there are only two dangers and there are cave-sets:

(1,2)
(2,3)
(1,3)
(1,3,4),

[(iv)] N is the total number of dangers of the type being considered.

then cave 4 can not possibly contain one of the given dangers. An algorithm to find all safe caves is to construct all combinations of N caves from the set of all caves that are members of any cave-set, eliminate combinations which do not fulfill the requirements (i.e. do not have at least one cave from every cave-set), and mark those caves which are not a member of any of the remaining combinations as safe. This algorithm is roughly as efficient as the first algorithm, but, unfortunately, it is sometimes not possible to explain the results of this last algorithm in simple terms. This is especially true in those unusual cases where the first algorithm failed. For example, the explanation for the example above would read, "*Cave 4 can not contain a pit because there are only two pits. If a pit were in cave 4, there would only be one other pit to explain all the evidence. We know that there must be a pit in either cave 2 or cave 3, but if the remaining pit were in cave 2 there could not be a pit in one of caves 1 and 3 (and we know that there must be a pit in one of caves 1 and 3). Likewise, if the remaining pit were in cave 3 there could not be a pit in either of caves 1 and 2 (and we know that there must be a pit in one of caves 1 and 2). Therefore there can not be a pit in cave 4.*" If this explanation seems rather obtuse, the reader is cautioned that the case explained above is the simplest possible case in which the first algorithm fails. Most other cases are significantly more complicated (and virtually unexplainable).(v) For this reason, the first algorithm was used instead of the second algorithm.

(v) Interested readers are challenged to write an explanation in simple terms for why cave 11 can not contain a pit if there are three pits and cave-sets (1,7,3,11), (1,2,3), (3,4,5), (5,6,7), (7,8,9), and (9,10,1). The explanation should not contain any new terms like "cave-sets".

The Probability Estimator

The third step of the Locator Algorithm estimates the probability that each cave contains a danger. It considers all the cave-sets, complete and partial, and is guided by five basic principles. The first principle (P11 of Figure 3.5) is that each member of a cave-set has an approximate probability of 1/N where N is the number of caves in the cave-set. In the case of partial cave-sets, N is an estimate of the size of the complete cave-set. For example, in Figure 3.4C the partial cave-set is caves 1 and 9, but N would be four as one would expect that the neighbors of cave 4 would add roughly two more caves to the cave-set. In any case, the probability of 1/N is not exact because it is often possible that some cave-sets will contain more than one dangerous cave. However, the estimate does give a relative measure of the risks involved with the different caves.

The next general principle is that if a cave has been identified as definitely containing a danger, that cave will explain the warnings responsible for any cave-sets of which it is a member. (vi) This makes it less likely that the other members of these cave-sets will contain a danger. The third principle is that if a cave is a member of more than one cave-set, it is more likely that the cave does, in fact, contain a danger. Conversely, the other members of these cave-sets are less likely to contain dangers. Lastly, it is better to shoot into caves that are likely to contain the Wumpus before visiting them. This principle considers the fact that if the Wumpus is not in the cave, then the arrow

---

(vi) Every cave-set with more than one member can be attributed to an originating cave at which a warning was received.

Figure 3.5

P11; Equal Likelihood Principle set $P=1/N$ where N is the number of caves in the smallest cave-set of which said cave is a member. It corresponds to the simplifying assumption that each cave-set had exactly one danger randomly assigned to it.

P12; Redundant Evidence Principle When it is noted that a given cave-set is a superset of another cave-set, then the superset cave-set is no longer considered as the warnings involved are completely explained by the subset cave-set. For the caves remaining in the superset after the subset is removed, the probability is set to:

$$P=\frac{U}{T-I-S}$$

where: I= the number of identified dangers. (The number of independent complete cave-sets).
U= the number of unidentified dangers. (The total number of dangers minus I).
S= the total number of safe caves based on the present evidence.
T= the total number of caves.

This probability is also applicable to caves for which there is no evidence to estimate the danger by.

P13; Multiple Evidence Principle for these caves which are members of more than one cave-set the probability is set to:

$$P_1 = 1 = 1 \times (1-\frac{1}{N_A}) \times (1-\frac{1}{N_B}) \times (1-\frac{1}{N_C}) \times \cdots$$

where: N= the number of caves in the given cave-set of which said cave is a member.

This expression corresponds to the simplifying assumption that exactly one danger is randomly assigned to each cave-set. It is P(danger)=1-P(no danger) where
P(no danger)=P(no danger from first cave-set)*P(no danger from second cave-set)*...

P14; Adjust For Multiple Evidence Principle for caves which are members of cave-sets to which the above formula is applied, the probabilities are adjusted to:

$$FP_I = 1 - \frac{1}{M_1}\left\{[(P_{1A}-\frac{1}{N_1}) + (P_{1B}-\frac{1}{N_1}) + \cdots] \cdot \frac{1}{(N_1-1)} + [(P_{2A}-\frac{1}{N_A}) + \cdots] \cdot \frac{1}{(N_2-1)} + \cdots\right\}$$

where: FP= the final probability for the specified cave.
N= the number of caves in the specified cave-set.
P= the probability initially calculated (above) for the specified cave.
M= the number of cave-sets of which the specified cave is a member.

Numeric subscripts are for the various cave-sets of which the Ith cave is a member.

Letter subscripts associated with numeric subscripts are for the various members of the specified cave-set, omitting the Ith cave.

This expression averages the changes from P13 out amongst the other members. It has no theoretical justification, but it works very nicely.

P15; Shooting Principle Whenever P_{Wumpus} exceeds $P_{Cut\ Off}$ it is reset to:

$$FP_W = (1 - P_W) \times \frac{1}{3} \quad \text{where:} \quad P_{Cut\ Off} = (1 - P_{Cut\ Off}) \times \frac{1}{3};\ P_{Cut\ Off} = .25$$

will start ricocheting with roughly a one third chance of killing the player. It is applied only if there are sufficient arrows so that the loss of an arrow will not end the game. These five principles are stated as rules with their accompanying formulas in Figure 3.5. The formulas are not referred to in any explanations but are simply used internally to allow the Advisor to evaluate the relative merits of the different moves more precisely.

Combination of the Results

To complete the comparisons between the different moves, the Expert calls the Locator Algorithm three times, once for each respective danger. Then it combines the probability estimates to come up with an overall estimation of the value of visiting each cave. The Expert is guided by eight basic principles in computing the overall evaluations. They are:

C0: Safe unvisited caves are preferable to all others.

C1: Given that it is necessary to take some risks, those moves with the lower risks for all dangers are preferable.

C2: Risking multiple dangers is worse than risking a single danger, presuming there is a common danger and that the common danger is equally likely in both cases.

C3: Presuming that there are common dangers that are equally likely, it is preferable to visit those caves where the unequal risks are lower.

C4: Bats are not as dangerous as pits or the Wumpus.

C5: Locating the Wumpus is an important goal, and it is worthwhile to visit caves which are likely to give information about the Wumpus' location.

C6: When the player is in a very dangerous section of the warren with no known route to a safer area, it is sometimes advisable to seek out bats in the hope that they will carry the player to a safer section of the warren.

C7: When the location of the Wumpus is probably known, it is best to shoot into the cave.

Figure 3.6
The Combination Principles were converted into the following formulas:

$$\mathrm{COST} = 1-(1-C_{Bat})*(1-C_{Pit})*(1-C_{Wumpus}) \qquad \text{(C0 thru C3)}$$

$$C_{Wumpus}=P_{Wumpus}$$

$$C_{Pit}=P_{Pit}$$

$$C_{Bat}=P_{Bat}*\frac{E_{Pit}+N_{Wumpus}}{N_{Caves}-E_{Bat}} \qquad \text{(C4)}$$

$$\mathrm{GAIN} = V_{Cave}+V_{Bat}$$

$$V_{Cave}=B*(1-P_{Bat})(1-P_{Pit})(1-P_{Wumpus}) \qquad \text{(C5)}$$

$$V_{Bat}=P_{Bat}*(1-P_{Pit})(1-P_{Wumpus})*\frac{N_{Caves}-N_{Visited}-E_{Pit}-N_{Wumpus}}{N_{Caves}-E_{Bat}} \qquad \text{(C6)}$$

where: P= The probability for the specified danger for the cave being considered.
B= A measure of the value of the information to be gained at that cave. (normally one). It is increased if the cave is likely to give information about the Wumpus.
N= The number of caves of the type specified.
E= An estimation of the number of caves with the specified danger that do not contain another danger of a higher priority. The priorities from highest to lowest are Wumpus, pits, and bats. For example,

$$E_{Bat}=N_{Bat}*(1-\frac{N_{Pit}}{N_{Caves}})(1-\frac{N_{Wumpus}}{N_{Caves}})$$

The above formulas can be used to give an index of:

$$\mathrm{INDEX}=\frac{\mathrm{GAIN}}{\mathrm{COST}}$$

The caves can then be ordered according to the INDEX. If COST=0, the INDEX would go to infinity; there would have to be different sizes of infinity. Instead, if the COST=0, the INDEX is set to very large number times the GAIN.

27

These principles were incorporated into the formulas shown in Figure 3.6 and are used to compute the index for each cave. This index is an estimation of the value of visiting each cave.

Summary

It is worth noting that the Expert described herein does not always find the best move possible. In several cases it was decided to use less accurate algorithms in order to make the computations of the Expert more understandable to humans. An example of this at a global level is the manner in which cave-sets are implemented. In Figure 3.7 there is a diagram of a section of the warren. The actual Locator Algorithm would conclude that caves 5, 6, and 7 are all equally likely to contain the Wumpus. In fact it is more likely that cave 6 contains the Wumpus, as a Wumpus in cave 5 or cave 7 does not explain all the warnings directly. (vii) An improved algorithm could apply probability rules throughout the propagation of cave-sets. A simplistic approach to this would be to assign to cave 3 and cave 4 a probability of one half of being next to the Wumpus. (viii) In this manner one achieves the results shown beneath Figure 3.7. These results are more realistic, but the improved algorithm has the unfortunate drawback of destroying the precise division between the probability rules and the logical rules. It also makes it much more difficult to break the rules into bite-sized chunks which are easy to teach. For such reasons, the improved Expert does not carry out its computations in the most accurate manner possible.

(vii) An additional connection is required between cave 5 and 7 to explain all of the warnings.

(viii) This is one of the general approaches suggested in Stansfield, Carr, and Goldstein (1976).

Figure 3.7

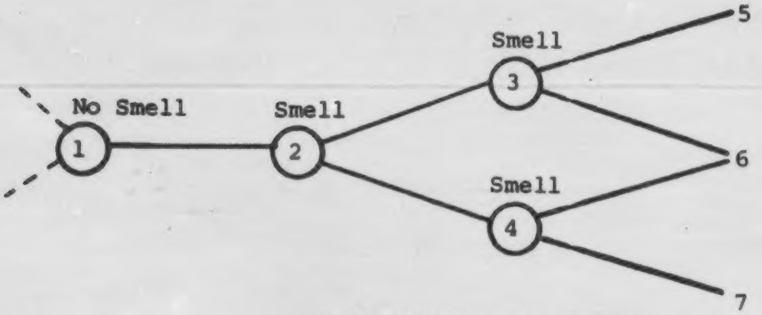

Probabilities For Wumpus

Cave	Old Algorithm Probability	New Algorithm Probability
5	.3333	.25
6	.3333	.50
7	.3333	.25

The sixteen danger specific rules (the eleven logical rules in Figure 3.2, and the five probability rules in Figure 3.5) together with the eight combination rules are a concise definition of the Wumpus Expert. Many of the intricacies of the expert, such as the priorities of the different probability rules when there are conflicts, have been omitted. Such details were worked out in the actual program, but are not worthy of note in this paper. The different rules have been presented in (and are numbered in) order of increasing complexity and, to a certain degree, in order of logical dependency. Normally, each rule uses only the input data and the results of lower numbered rules. The 24 rules taken together comprise the knowledge which the Wumpus Advisor endeavors to teach. In this sense, the entire Wumpus Advisor centers around the Expert.

The Improved Wumpus Advisor, Wusor II

Chapter 4

The Overall Structure

Wusor II has been implemented with the structure shown in Figure 4.1. The Wumpus Expert module is an implementation of the expert just described. The Move Comparer analyzes the different moves noting the knowledge required to justify their relative merits. The Psychologist uses this information to determine how the player's actions (moves) reflect upon his knowledge of the game and informs the Student Model as such. The Student Model can then decide what the player is presumed to know, and, from this, it determines what material is at the student's level. (i) The Wumpus Advisor uses this information (added to the results of the Move Comparer) to determine what to explain to the student.(ii) If it decides to give advice to the student, it will call upon the English routines to explain certain of the results of the Expert. The English routines use the Student Model to prune their explanations in light of the player's current knowledge. The Wumpus Executive coordinates the activities of the Wumpus Advisor and Wumpus game and implements the conceptual model of a tutor shown in Figure 1.1.

(i) Note that within this paper, "Student Model" will be used to refer to more than just the data about the student, but also the routines which maintain and evaluate the data. In effect, the Student Model is a "black box" which answers queries about the student's current state, and we will not be concerned about the details of its functioning until later.

(ii) Burton and Brown (1975) describes a previous Advisor program for the game of West which was built around an Expert and which utilized a primitive Student Model.

As the Wumpus Expert performs its calculations, it keeps a record of the processes it went through to arrive at its results. The information relevant to each cave is kept on the cave's property list so that the information is easily accessible. Tags are also put on the property lists which correspond to the unique numbers assigned to each of the rules listed in Figure 3.2 and Figure 3.5. In this manner the Expert provides a complete set of sub-conclusions and sub-results along with the final results. These property lists are used extensively by the English Generator and Move Comparer. Both of these modules contain a structure modelling that of the Expert.

Figure 4.1
Diagram of Wusor II
Arrows Represent Data

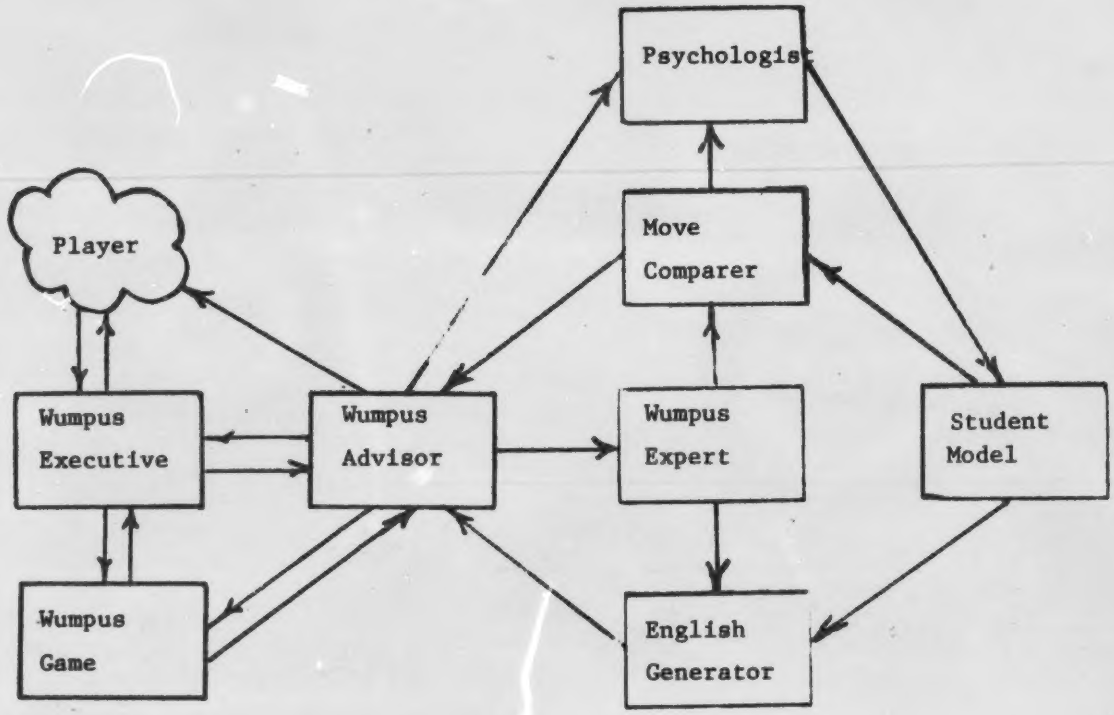
Psychologist
Player
Move
Comparer
Wumpus
Executive
Wumpus
Advisor
Wumpus
Expert
Student
Model
Wumpus
Game
English
Generator

The Wumpus Advisor also has the ability to modify the Wumpus game to create situations more conducive to learning. This has been done primarily by allowing the Wumpus Advisor to define the number of bats, pits, Wumpii, and caves appropriately as well as choosing the initial caves encountered by the student. The starting cave is selected by the Advisor, and the player's first move is transposed to the desired cave (all of the choices for the first move are isomorphic). This prevents the player from losing on the first move before he has sufficient information to make intelligent decisions and also allows the Wumpus Advisor to challenge the student with simple idealized problems that are appropriate for a student of his level. Examples of problems selected for students of differing levels are shown in the Appendices. As a further addition, the Wumpus Advisor could be given the ability to restructure the warren as the game progresses. This modification is non-trivial since these changes must be transparent to the player, but, more importantly, the Wumpus Advisor must be able to fully describe the situation it is trying to create. The current Wumpus Advisor does not have sufficient information available to it to allow it to fully specify more complicated situations, but it could attempt to move dangers so that, if the student were playing wisely, he would gain positive reinforcement. Conversely, if the player were playing unwisely, negative reinforcement would be supplied.

In the remainder of this chapter we will discuss the curriculum of the Wumpus Advisor. In the next chapter, we will fully describe the Student Model and its relationship to the Move Comparer and the Psychologist. Then in Chapter 6 we will discuss the tutoring strategies which have been selected for Wusor II. In Chapter 7 we will describe the English Generator before we discuss (in Chapter 8) the experimentation which is appropriate for the Wumpus Advisor.

The Curriculum

There is a definite hierarchy for the different rules used by the Expert, which makes it logical to have a corresponding hierarchy in the curriculum of the Wumpus Advisor. It is not realistic to explain the different probability rules until the student has mastered some of the logical rules used in the derivation of the cave-sets because the probability rules are so dependent on these cave-sets. Because of such dependencies, the curriculum of the Wumpus Advisor has been broken down into five phases.(iii) In each phase, the Wumpus Advisor attempts to teach certain rules, and the student will not be advanced to the next phase until he has mastered those rules which are the foundations for the next phase. In the zeroth phase, the student is taught the basic rules of the game. In particular, the player is taught that he is allowed to backtrack, (iv) and that, if a cave neighbors on a cave in which there were no warnings, that cave is necessarily safe. In phases one through four, the player is allowed to advance separately with respect to each danger and each danger has its own curriculum. This is necessary because the complexity of the different rules varies widely according to the distance through which warnings propagate. For this reason, there are two basic curriculums: a curriculum for when warnings propagate one cave and a curriculum for when warnings propagate further.

(iii) BIP, (Barr, Beard, and Atkinson, 1975) a CAI program for teaching elementary programming skills, has considered the issue of the conceptual dependencies between the skills which are taught. BIP takes the more general approach of formally stating these dependencies in an Information Network. This task is performed implicitly in the phases of the Wumpus Advisor, along with various other functions.

(iv) Knowledge of backtracking is essential for good play, and some new players do not grasp it on their own.

The curriculums for bats and pits presumes a warning distance of one cave. In the first phase of this curriculum, the player is taught how to perceive which caves are safe from the absence of a warning. He must master these rules before he can advance to phase two where he is taught the first probability rule (P11). The player must master this probability rule and its supporting logical rules before he can advance to phase three. In phase three, the player [illegible] taught the reasonably complicated probability rules (P12 and P13) and must master both of them before he can advance to the last phase. In the last phase, the player is taught the most advanced of the probability rules (P14). This rule is left to the last as it has a conceptual dependence on the two preceding rules. All of the rules taught in each phase are shown in Figure 4.2 along with the necessary conditions for advancement to the next phase. (v)

(v) In Figure 4.2, there is a new rule, T19, in the first phase of the Wumpus curriculum. This rule is a specialization of the logical rule L5. It is relevant whenever a "More Than" classification involves N applications of rule L5 where N is the warning distance. It is necessary in cases where the warning distance is greater than one because in these cases students will sometimes not apply L5 to the limit.

Figure 4.2
Rules taught in each phase.

Phase	0	1	2	3	4
Bats	L1	L0 L3 L5	L4 L7 P11	L6 L10 P12 P13	P14
Pits	L1	L0 L3 L5	L4 L7 P11	L6 L9 L10 P12 P13	P14
Wumpus	L1	L0 L2 L3 L5 T19	L4 L7	L6 L8 L9 L10 P11 P15	P12 P13 P14

Rules which must be mastered to advance to the next phase.

Phase	0	1	2	3
Bats	L1	L3 L5	L4 L7 P11	P12 P13
Pits	L1	L3 L5	L4 L7 P11	P12 P13
Wumpus	L1	L3 L5 T19	L4 L7	P11

The Wumpus curriculum presumes a warning distance of more than one cave. In the first phase of the Wumpus curriculum, the player is taught the logical rules necessary for perceiving which caves are safe from the Wumpus. This phase has the same rules as in the bat/pit curriculum, but because of the increased warning distance, the rules are more complicated to apply. In the second phase, the Wumpus Advisor teaches the player those rules necessary to arrive at cave-sets. This separation of the rules for deriving cave-sets from the first probability rule is made possible because of the special value assigned to caves which are likely to give information about the Wumpus. The Advisor can speak of the value of completing cave-sets without actually getting involved with the details of probability. The separation of the logical rules used in the derivation of cave-sets from the probability rules is necessary because of the extra complexity of a warning distance of two caves. This not only makes the logical rules more complicated, but also makes the probability rule P11 more complicated as it now involves estimates for the size of the cave-set. In the third phase, the player is taught P11 and P15. He must have mastered P11 before he can go on to the fourth and final phase. In the last phase the Wumpus Advisor attempts to teach the P12 and P13. This is not possible in the standard game as there is only one Wumpus and there are no situations to demonstrate these rules. As a result, the Wumpus Advisor changes the game so that there are two Wumpii in the warren as soon as the student advances to the last phase. (vi)

The player is allowed to advance independently through the last four phases of each danger. This allows for the possibility that the player might learn the use of a rule with respect to one danger, but might not recognize that it also applies to another danger. However, once the player has mastered a few of the logical rules for bats and pits, it can be expected that he will recognize their correspondence for other rules. (vii) For this reason, there is a function in the Student Learning Model which notes when the player has learned several corresponding rules for bats and pits; thereafter, it presumes that the player will recognize the correspondence between any other rules for bats

---

(vi) Note that since there is only one Wumpus for the earlier phases, the logical rule which considers the number of dangers (L10) becomes much simpler and is taught in the same phase with P11.

(vii) This statement is based upon the players which the author has observed. The author has observed fifteen or so students in their interactions with Wusor; in all cases the students quickly learned the correspondence between applicable bat and pit rules. However, this issue warrants further experimentation.

and pits and adjusts the *Familiarity Values* accordingly. The player is taught the combination rules (rules C0 thru C7) independently of the different phases. However, in teaching the Combination rules, the Wumpus Advisor takes into consideration the conceptual dependencies of the rules.

The curriculum that is taught by the Wumpus Advisor is designed to first teach the student the basic rules and then the more complicated applications of these rules. The total knowledge is broken down into specific items, and these items are grouped into phases so that each phase has applications that can be taught directly. This allows the Wumpus Advisor to make short but meaningful explanations to the student at all times.

The Psychologist and Student Model

Chapter 5

An Overview of the Move Comparer, Psychologist, and Student Model

The Student Model along with the Move Comparer and Psychologist have the purpose of determining just what moves can be explained to the student and of determining what rules can be pruned from the explanations to be given. These two goals require that the Student Model be able to determine which rules are known to the student. If a human tutor were faced with such a problem, his Student Model would be based on several sources of information. This sort of generalized Student Model is shown in Figure 5.1. The decision of whether or not the student knows a particular rule is based on the following inputs.

1. How often the student seems to have applied the rule.
2. How often the student has had the rule explained to him.
3. The student's knowledge of other rules (upon which this rule depends).

Consideration must also be given to other important points. They are:

4. How quickly the student learns.
5. The student's initial knowledge of the rules.
6. The time period that has elapsed and how forgetful the student is.

Of course, these are not the only things considered by a human tutor but they are certainly very important considerations. Our Student Model will take all of these inputs into consideration.

39

Figure 5.1
Generalized Student Model

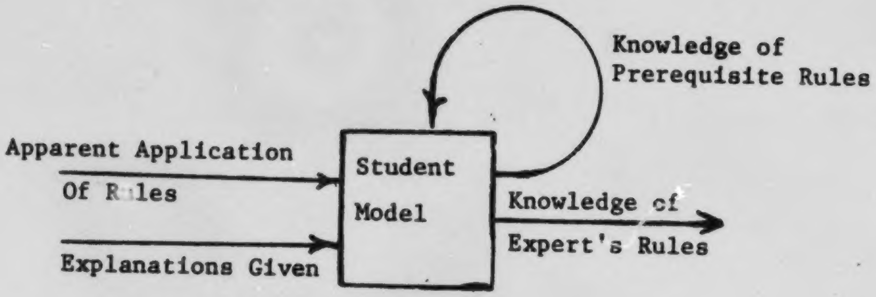

In Figure 5.2 there is a diagram of the data flow between the Move Comparer, Psychologist, and Student Model as well as the various submodules of the Student Model. The Move Comparer takes any two moves and determines which rules are necessary to justify their relative merits with respect to a single danger.(i) It also queries the Student Model to determine whether or not these rules are known to the student, acceptable for teaching, etc.. The Psychologist uses the results of the Move Comparer according to the Combination Rules to determine if the student has demonstrated a working knowledge of any danger-specific rules while the Wumpus Advisor uses them according to the Combination Rules to determine which of the Expert's results it will explain to the user. The loop in Figure 5.1 representing the student's knowledge of other rules is implemented through the loop from the Learning Model to the Move Comparer to the Psychologist to the Knowledge Model and back to the Learning Model. The Student Knowledge Model maintains ***Familiarity Values*** which represent the student's familiarity with any given rule. The Student Learning Model uses its estimate of the student's learning ability to interpret the ***Familiarity Values*** and determine which rules the player is presumed to know. (ii) With these rules, the Student Learning Model can implement the curriculum of the Wumpus Advisor by computing which rules are acceptable for teaching. Finally, the Student Model Critic analyzes the interactions between the Psychologist and the Learning Model and adjusts the Learning Model as appropriate. (iii)

(i) In Appendix C there is a complete description of how the Move Comparer functions.

Figure 5.2
Arrows Represent Data

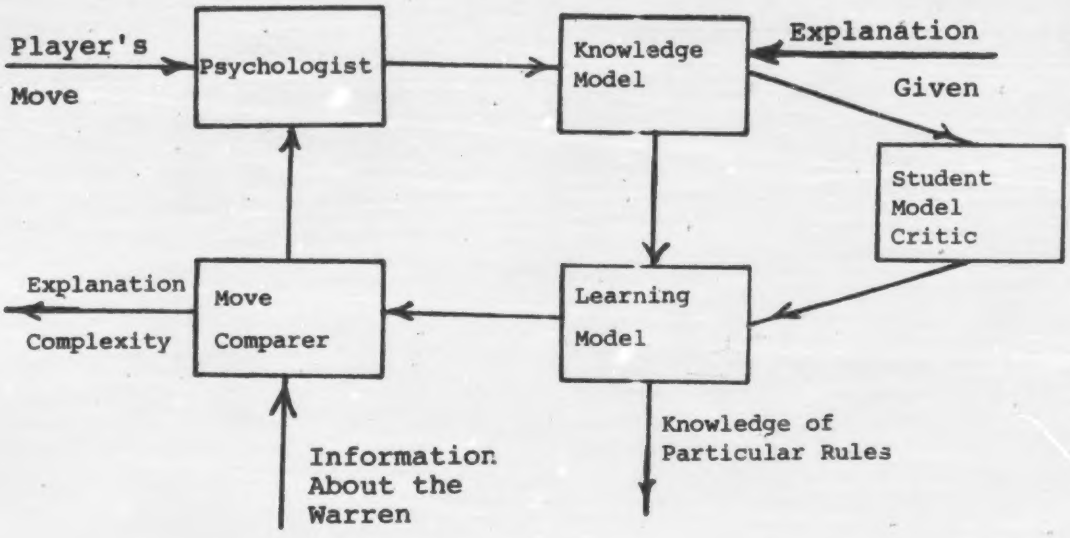
Player's Move
Psychologist
Knowledge Model
Explanation Given
Student Model Critic
Explanation Complexity
Move Comparer
Learning Model
Information About the Warren
Knowledge of Particular Rules

(ii) In this chapter there will be many references to the student's knowledge of the different rules. This will normally refer to the Student Learning Model's estimate of the student's knowledge.

(iii) The Student Model also maintains disc files on all of the users. These files are keyed on the student's first and last name; any time a new player uses the Wumpus Advisor, a new file is created. It is copied onto disc at the end of each session, and his file is reloaded at the start of a new session. The file contains the variables and arrays of the Student Knowledge and Learning Models as well as various other information.

The Psychologist and Combination Rules

The Psychologist begins its analysis when the player asks to visit a cave in the fringe area (iv), and it assumes that, if the player has moved to a cave which is distant from the starting cave, he must have discerned that something was wrong with those moves which were closer to the starting cave. (v) The assumption that the player recognized a fault with each of the worse moves is qualified by considering the route which the player took to his move. If the route taken was non-optimal, the Psychologist will consider only those caves which were very close to the starting cave. (vi) If the player takes a particularly lengthy route, the Psychologist will consider only the immediate neighbors of the starting cave. (vii) Likewise, if the Wumpus Advisor has recently advised the player to go to the move which he finally selected, the Psychologist will not consider his move as it can not be certain just how much of the

---

(iv) The fringe area contains those caves which the player has not visited and which are accessible to the player through caves which he has already visited.

(v) The starting cave is the last fringe cave which the player visited or the last cave at which the player shot an arrow or received advice from the Wumpus Advisor.

(vi) The Wumpus Advisor has several features which are intended to aid the player during the game. One of these is a route planner which will explain an optimal route to any cave that is accessible through visited caves. If the player uses this feature to find a route to a cave and then goes there directly, the Psychologist will presume that the player recognized some fault in *all* caves which were worse than the selected move.

(vii) Random chance becomes very likely when the player takes an extremely indirect route. Therefore, the Psychologist only considers caves that are within N caves of the starting cave where $N=2O-L-1$, L is the length of the route actually taken by the player, O is the length of the optimal route, and N has a minimum value of one.

advice the player actually understood. (viii)

Presuming that there were worse moves for the Psychologist to consider, it would determine if the player has demonstrated a working knowledge of any rules. The Psychologist would conclude that the player has demonstrated a working knowledge of a group of rules with respect to a certain danger only if there were no other plausible explanation for why the player's move was the better move. These conclusions are based on a restrictive definition of the first four Combination Rules. There are examples of this in Figure 5.3. In Figure 5.3A, the last fringe cave visited by the player was cave 5. The player then moved directly to cave 12. Cave 6 is a worse move because it risks pits and the Wumpus. However, in this case the Psychologist would do nothing as it is not clear whether the player perceived the danger from pits and/or the danger from the Wumpus. The player also bypassed cave 3, a cave that was unsafe for bats, but, in this case, the Psychologist would presume that the player had perceived that cave 3 was unsafe (the default) and that cave 12 was safe from bats. It would note that the player had demonstrated a working knowledge of those rules which made cave 12 safe from bats. (Logical Rules L3 and L5 for bats and Combination Rule C0). In Figure 5.3B, the player started at cave 18 and chose the better move, cave 14, bypassing cave 17. In this case, there was a common danger from bats and a lesser risk from pits. The Psychologist would assume that the player saw that there was a smaller risk from pits in cave 14 because there was absolutely no difference in bats for the player to perceive. It would

---

(viii) Correspondingly, the Psychologist will not consider worse moves which the Advisor had previously informed the player were bad.

note that the player had demonstrated a working knowledge of the rules necessary to explain both caves' probabilities for pits. (Logical Rules L1, L4, and L7 for pits, Probability Rules P11 and P12 for pits, and Combination Rule C3). (ix)

Figure 5.3

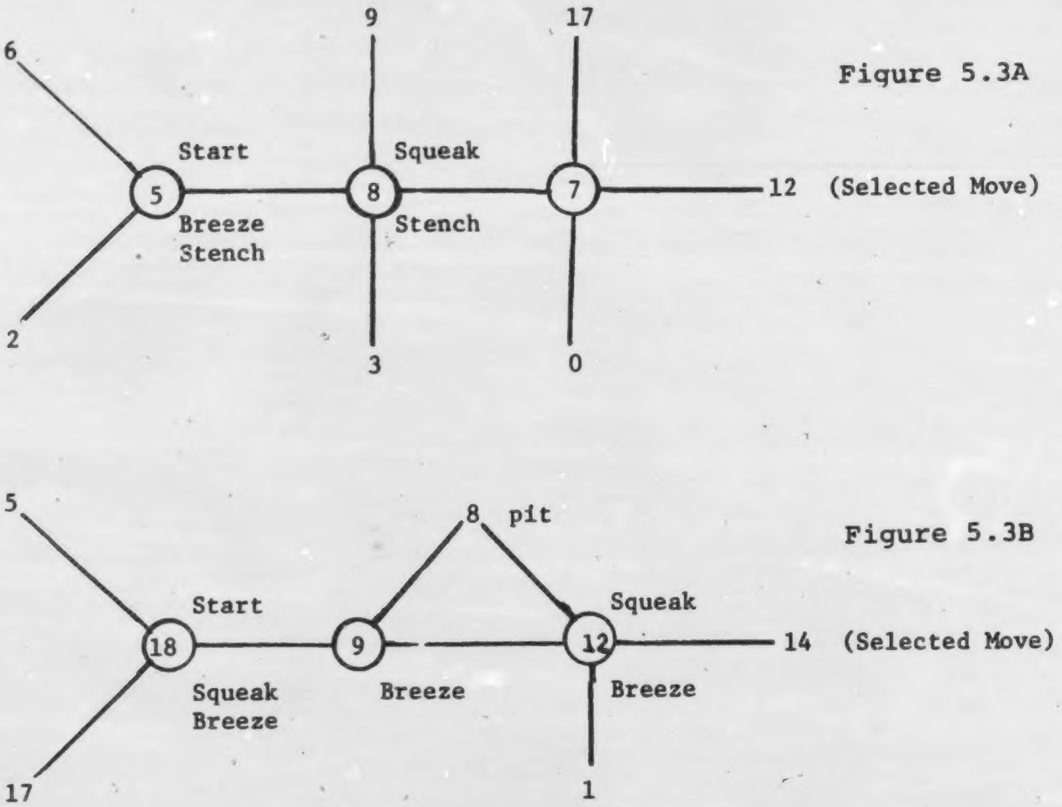
Figure 5.3A
6
9
17
Start
Squeak
5
8
7
12
(Selected Move)
Breeze
Stench
Stench
2
3
0
Figure 5.3B
5
8
pit
Start
Squeak
18
9
12
14
(Selected Move)
Squeak
Breeze
Breeze
Breeze
17
1

(ix) The Wumpus Advisor also implements the different Combination Rules, but it uses a less restrictive definition of the Combination Rules. While the Wumpus Advisor will only explain why a move is better because of a single danger, it is acceptable for the move to be better for reasons other than those explained. If the player asked to move to cave 6 in Figure 5.3A, the Wumpus Advisor would advise him to move to cave 12 instead because of a smaller risk from the Wumpus. The Wumpus Advisor would only cursorily mention the smaller danger from pits if it mentioned pits at all because it tries to keep its explanations short.

Sometimes a player will make a lucky move which gives the appearance of a mastery of certain rules but which does not really indicate a knowledge of said rules. To restrict this possibility, the Psychologist first checks to insure that the rules in question are not well above the expected level of the student.(x) If they are, the Psychologist will consider his move a lucky move. As a further precaution, the Psychologist compares the move selected by the player with all the better moves available to him. If there are any moves which are better than the selected move and which the user should have known were better, the Psychologist will degrade the Student Knowledge Model appropriately. The Psychologist determines if the student should have recognized the better move by insuring that the move in question was better than or exactly the same as the selected move in every way. It then insures that the student was thought to know all those rules which indicated the quality of the better move, and, if so, it will degrade the appropriate rules of the Student Knowledge Model. In some cases, this degrading of the Student Knowledge Model will cause the Student Learning Model to decide that the player does not actually know the rule(s) in question, and it will move the player to as low a phase as necessary.

(x) Rules are considered to be above the expected level of the student if they are not rules which are considered acceptable for teaching.

The Student Learning Model

Just as the Student Knowledge Model can be viewed as an overlay of the rules (and knowledge) of the Expert, the Student Learning Model is an overlay of those tutoring capabilities possessed by the Advisor. It is composed of those abilities which are deemed appropriate for each student.(xi) In Goldstein's Coach proposal (1977), an Advisor is envisioned capable of different explanatory strategies so that the Learning Model contains those explanatory strategies which are appropriate for the particular student. However, the Wumpus Advisor does not have this degree of freedom (as it has only one explanatory strategy)(xii), and the Learning Model is primarily concerned with a vector which represents each student's learning ability in the three dimensional *Learning Space*. The first dimension of the *Learning Space* is how long it takes the student to learn something (the *Repetition Factor*) and the next dimension is how long it takes him to forget something (the *Forgetfulness Factor*). The *Receptivity Factor* is the last dimension and indicates how often the student likes to be spoken to. How often the Advisor gives the student advice depends on the student's *Receptivity Factor*. There is a large region within the three dimensional *Learning Space* in which the Wumpus Advisor can function effectively.

The first dimension of the student's learning ability (the *Repetition Factor*) is a variable which determines how high a *Familiarity Value* the student must have before it can be presumed that he knows a

---

(xi) Overlay modelling is discussed more fully in Carr & Goldstein (1977).

(xii) Later versions of the Wumpus Advisor will have expanded explanatory abilities with a correspondingly increased Learning Model.

rule. The *Familiarity Value* for a rule is incremented for one of two reasons. If the player is not deemed to have *heard of* the rule in question, then the *Familiarity Value* will be incremented when the player is told the rule in question.(xiii) Likewise, the *Familiarity Value* will be incremented if the student demonstrates a working knowledge of the rule (as determined by the Psychologist). (xiv) Whenever the *Familiarity Value* exceeds the *Repetition Factor*, the Student Learning Model will presume that the student knows the rule in question. (xv) When it notes that the student has learned a rule, it also checks to see if he has learned the rules requisite to advancing to the next phase. If so, the Student Learning Model will note that it is acceptable to teach the additional rules of the new phase. (xvi)

The second dimension of the student's learning ability is the *Forgetfulness Factor*. The Student Learning Model normally presumes that the student has forgotten certain of the more recently acquired rules between sessions. This is done by decrementing the *Familiarity Values* according to the *Forgetfulness Factor* and how long the student has been

---

(xiii) Whether or not the player is deemed to have *heard of* the rule in question is decided by comparing the number of times he has been told the rule with a variable that is dependent on his *Repetition Factor*. Note that this requires the Student Knowledge Model to keep track of how often the player has been told a rule as well as his *familiarity*.

(xiv) *Familiarity Values* are a rather primitive implementation of the Model in Figure 5.1 as the two external inputs are combined linearly.

(xv) It is worth noting that the test value for having *heard of* rules is such that it is not possible for the Student Learning Model to presume that the player knows a rule without the player having demonstrated a working knowledge of the rule at some time.

(xvi) Even after the student is presumed to have mastered a rule, the *Familiarity Value* is updated whenever the student demonstrates a working knowledge of the rule.

away from the Wumpus Advisor. [xvii] The *Familiarity Values* are decremented by an amount equal to the *log* of the number of days that the player has been away from the Wumpus Advisor (plus one). The base of the *log* is the inverse of the *Forgetfulness Factor.*

The Student Model Critic

The Student Model Critic analyzes the player's behavior whenever the player makes a move which indicates a working knowledge of a rule which he is not presumed to know, and when he makes a move which he should have recognized as better (these situations are identified by the Psychologist, of course). In such cases the Critic will often adjust the student's *Repetition Factor* or *Forgetfulness Factor*, whichever is appropriate. Also, if the player follows the advice of the Wumpus Advisor, the Critic will increase his *Receptivity Factor*. In contrast, if the player ignores the advice of the Advisor and makes moves which the Advisor has noted as bad, the Critic will decrease the player's *Receptivity Factor.* [xviii] These small incremental changes to the key variables of the Student Learning Model are intended to adjust the Learning Model to the actual student's learning ability.

The Student Model Critic makes an analysis of the situation before making any changes to the *Forgetfulness Factor* or *Repetition Factor*. Its analysis is guided by the six Critic rules shown in Figure 5.4. The Critic first notes whether or not the player is a new player; in such cases, it is quite likely that the initial estimate of the student's knowledge was incorrect, and so the Critic will only make immediate adjustments to the Knowledge Model. This is formally stated in the first two Critic rules of Figure 5.4. The Critic will presume that it has an acceptable approximation in the Knowledge Model when it has to undo a decision made earlier. For example, if the Critic repeatedly decreased its estimation of the student's knowledge, it would stop adjusting the Knowledge Model as soon as the player learned a rule which he had previously demonstrated that he did not know.(xix) This is formally stated as the third Critic rule. Once the Critic feels that the Knowledge Model has been adequately adjusted, it will begin adjusting the ***Repetition Factor*** and ***Forgetfulness Factor***. However, the Advisor will not make any changes to the Factors if the student's ***Familiarity Value*** for a rule in question is not relatively close to the ***Repetition Factor***. If it is not fairly close, it is far more likely that this case is not indicative of the student's learning ability, but, instead, a fluke (a mistake on the part of the player or a lucky move).(xx)

[xvii] The values representing how often the player has been told each rule are also decremented as his memory of the advice given is presumed to have decreased over the given time period.

[xviii] The amount which the Critic increases or decreases the student's receptivity depends on how recently the advice was given.

Once it is determined that the move in question involves a rule with a marginal ***Familiarity Value***, the Critic will make the adjustment which is indicated by the last two Critic rules.

(xix) Because game situations often allow the student to more quickly demonstrate a knowledge (or the lack thereof) of rules with respect to certain dangers, the Critic notes whether or not the Knowledge Model has been adjusted with respect to each danger.

(xx) The Critic will not make any changes to the Learning Model if the ***Familiarity Value*** was set in the initialization period.

Figure 5.4

CR1 If a new player demonstrates a working knowledge of a rule which he was not initially presumed to know, adjust the ***Familiarity Value*** to indicate a knowledge of said rule.

CR2 If a new player indicates that he does not know a rule which he was initially presumed to know, adjust the ***Familiarity Value*** to indicate that he does not really know said rule.

CR3 Cease adjusting the initializations of the ***Familiarity Values*** as soon as some decision made according to CR1 or CR2 is contradicted.

CR4 If the player demonstrates a working knowledge of a rule which he was not presumed to know, then

  A Decrease the player's ***Forgetfulness Factor*** if this was a rule which he was presumed to have forgotten recently.

  B Otherwise, decrease the player's ***Repetition Factor***.

CR6 If the player's actions indicate that he does not really know a rule which he was presumed to know, then

  A Increase the ***Forgetfulness Factor*** if this was a rule which the player was presumed to have learned in a previous session.

  B Otherwise, increase the player's ***Repetition Factor***.

In this manner, the Student Model Critic adjusts the values of the Student Learning Model as appropriate.(xxi)

The Student Model Critic also identifies those situations in which the limits of the Wumpus Advisor are exceeded and a human teacher is necessary. This is done by noting when the player does not seem to be advancing after an extended period and, more importantly, by noting those situations in which the limitations of the Wumpus Advisor are exceeded. For example, within a two dimensional *Learning Space* composed of the *Repetition Factor* and the *Forgetfulness Factor*, the Wumpus Advisor is

---

(xxi) At this point it should have become clear to the reader that the Student Model keeps track of when a player is presumed to have learned/unlearned each rule, which rules were forgotten, etc..

competent at advising students within the area shown graphically in Figure 5.5. The Student Learning Model can represent a user as a vector anywhere within the marked area.(xxii) Any time the student's Learning Vector falls outside of this region, the Student Model Critic calls for human help as the limits of the Wumpus Advisor will have been exceeded.

Figure 5.5

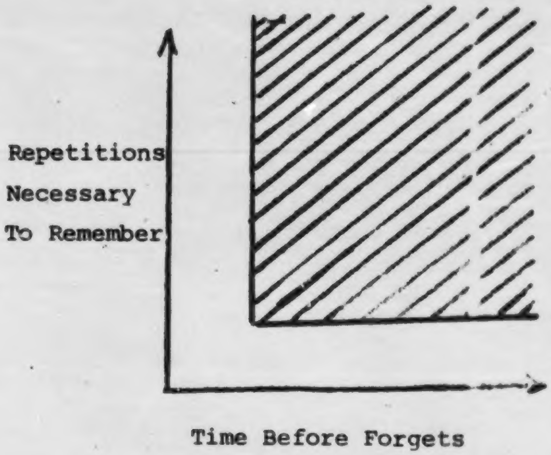

Because of the limitation on the ***Repetition Factor***, the Wumpus Advisor can not always keep up with extremely quick students, but the Advisor is not really necessary in such cases. However, the limitation of the ***Forgetfulness Factor*** may be very significant, as the Wumpus Advisor can not deal effectively with students who are extremely forgetful. The limitation is caused by the fact that the Wumpus Advisor keeps track of time in days, which could be too large a unit of measure in extreme cases. Further testing is required to determine if this is an acceptable limitation. Otherwise, it would be necessary to modify the Wumpus Advisor to keep track of time in hours or, possibly, minutes.

(xxii) The reader should note that the Wumpus Advisor is not able to finely tune its advice for players whose vectors are close to the limits of the Wumpus Advisor's abilities.

It is essential that the Student Model have the ability to recognize its own limitations. From observing various students and their interactions with the Wumpus Advisor, the author has noted that one very important step in the student's learning is the development of a good representation for information about the warren. Without a good representation, it does not seem to be possible to advance to the more advanced phases. In all cases, the students have eventually developed a good representation on their own (the Wumpus Advisor does not currently have the ability to teach good ways to draw warrens). However, it is likely that some students would not overcome this obstacle on their own. While it is simple to add in the ability to graphically display "good warrens" to the student, this is almost certainly only one of many possible pitfalls for the student that the Wumpus Advisor would not be able to deal with. In such cases, it is essential that the Wumpus Advisor have the ability to monitor itself and determine when something has gone awry.

The Initializations

The Student Model is initialized according to information elicited from the user. Whenever a user starts up Wusor and the Wumpus Advisor does not already have a file on disc about him, the Advisor will ask the new user a series of questions which are used to initialize the Student Model. First, the Wumpus Advisor will ask the student 1) how old he is, 2) how many years of education he has completed, and his general attitude to Math/Sciences. With this information, the Advisor will calculate the student's Repetition Factor and Forgetfulness Factor under the assumption that older, better educated students who enjoy the Math/Sciences are better prepared to learn about the Wumpus domain. Once the Student Learning Model has been calculated the Advisor will ask the student how many games of Wumpus (if any) he has played. This will be combined with whether or not the student chose to read the instructions and his estimated learning ability to select one of the four possible starting states for the Student Knowledge Model. From experimentation it has been found that the Advisor advances the Student Model more quickly and accurately than it degrades the Student Model (as increasing knowledge of the game is normal), and so the initial state of the Knowledge Model is generally a conservative estimate. In this manner the Student Model is initialized to a state which is hoped to be close to the actual state of the student. In any case, it is not of great importance that the Student Model be initialized with extreme accuracy as the Student Model has the ability to compensate for such errors.

Tutoring Strategies

Chapter 6

Model of Tutor/Student Interactions

In Figure 6.1 there is a model of the interactions between a tutor and a student. It is intended to emphasize the cyclical natures of their communications. The tutor communicates his explanations to the student and the student communicates his understanding to the tutor (feedback). The tutor modifies his instructions according to the feedback and will elicit feedback if he is receiving insufficient feedback. Typically, the instructions are the explanations of a teacher, while the feedback includes everything from facial expressions (boredom, non-comprehension, etc.) to test scores. This model is of a generalized tutor which could be a human instructor or a computer program.

Figure 6.1
Model of Tutor/Student Interactions

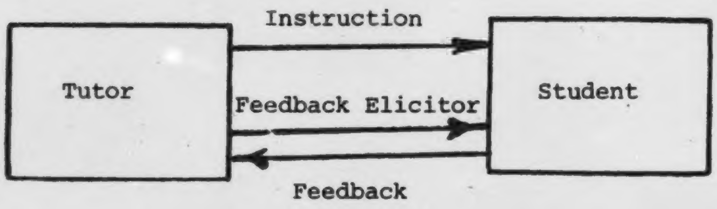

There are several limitations in the computer's ability to communicate. This difference is primarily due to the fact that the computer's communications are restricted to the domain of typed responses. The computer can not rely on facial expressions or verbal signals to emphasize key points. Also, it can not use facial expressions as a source of feedback. For example, it is inappropriate for a computer program to use jokes as a feedback elicitor because the computer can not evaluate the response (laughter, snoring, etc.).

In general, a computer program must rely on more formal avenues of communication for its instructions, feedback elicitors, and sources of feedback.

Instruction

The instructions can have various forms. These different approaches to teaching can be laid out along a spectrum which represents the amount of thinking required of the student. This is shown in Figure 6.2.(i) A point at the left end of the spectrum indicates that the tutor simply presents algorithms about the domain. This approach requires little thought on the part of the student and can be highly effective at quickly giving the student mastery of the field.(ii) The opposite extreme is to simply present the problems to the student and state the solutions. This requires the student to work out a technique for solving the problems. This process is generally very slow and painstaking and leaves open the possibility that the student will ***never*** master the given problem domain. However, once the student has mastered the domain with the latter approach, he is very unlikely to forget the learned material and those lessons which the student learns in the given domain are more likely to be applied in other domains.

(i) The idea qualifying the type of teaching according to the effort required of the student is presented in Winston (1970).

(ii) Of course this discussion presumes that the strategy selected will be developed in the best means possible. The approach of presenting algorithms to the student could be complicated by using very complex language, thereby requiring significant effort of the student just to comprehend the tutor. This is contrary to the intent of the algorithmic approach, and we do not consider such cases here.

Figure 6.2
Spectrum of Instruction Emphasis

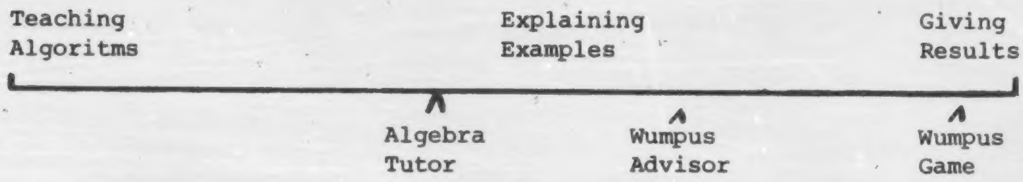

In developing the Wumpus Advisor, it was decided that the Wumpus Advisor's instructions should be biased toward explaining results. This is shown in the spectrum of Figure 6.2. The Wumpus game itself could be located at the extreme end of the spectrum (as shown in Figure 6.2) as it simply supplies the player with results and the player is left on his own to develop better techniques for avoiding dangers. The disadvantage of the game without an Advisor is, as stated earlier, that the player can reach plateaus in his development and not advance to higher levels. The Wumpus Advisor generates explanations which are intended to lead the student to the accepted solutions. With this purpose, the most direct approach would be to present the student with a somewhat simplified version of the algorithms used by the Expert, but this approach does not seem to be conducive to teaching better ways to think about problems.

The author has tutored various students in the game of Wumpus, and, when he presented the algorithms of the expert, the players quickly mastered the game, but there were also no indications that these students developed any improved thinking habits. However, when the author relied on the Wumpus Advisor and its explanations of the results (as opposed to the actual algorithm), he found that the learning process varied. Those players who were already logically oriented (such as M.I.T. graduate students) quickly developed algorithms of their own. The players who did not possess the desired thinking habits (i.e. those who were not mathematically/scientifically inclined) found the ideas presented intriguing and rather difficult. They would put serious thought into the explanations and after many tries mastered the material. Comments were made such as "*Oh wow, this is a completely new way of thinking*.", but after having mastered the material they found it difficult to accept that the material was so difficult ("*but its so obvious*").(iii) Because of such evidence, the Wumpus Advisor generally does not state its rules formally to the student but rather gives explanations of the results.

In domains where the material to be taught is thought to have more intrinsic value (such as Algebra which, it can be argued, can be of use in every day chores such as balancing checkbooks), there is a better justification for teaching with a more algorithmic approach. A computer tutor for Algebra might factor several problems, showing all the steps to the student. Then it could give the student various problems to do on his own to help him remember the algorithms presented. This form of instruction is appropriate in those domains where the subject matter has *intrinsic* value (rather than helping to develop improved thinking habits).(iv)

Feedback

It is also possible to lay out the different approaches to feedback

---

(iii) The above observations were made with about fifteen students in separate sessions, each of which lasted an average of three hours. All results were the subjective evaluations of the author.

(iv) Socratic instruction (through leading questions) can be viewed as another approach to the problem of non-algorithmic tutoring strategies. Socratic questioning is discussed fully in Collins.

along a spectrum. Such a spectrum is shown in Figure 6.3. The point at the left end of the spectrum is for tutors which ask the student direct questions about the subject matter. The point at the right end indicates that the tutor estimates the student's knowledge solely by observing his behavior. An Algebra tutor which presents problems to the student and observes how he solves them, would be toward the middle of the spectrum.

Figure 6.3
Methods of Gaining Feedback

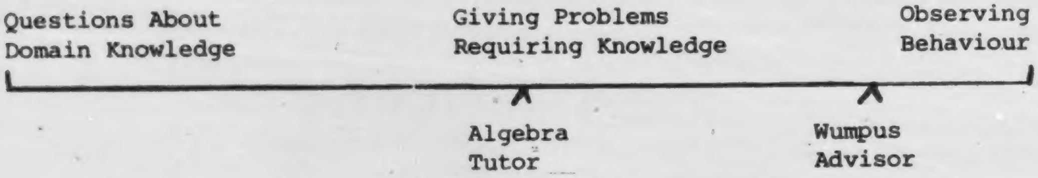
Questions About Domain Knowledge
Giving Problems Requiring Knowledge
Observing Behaviour
Algebra Tutor
Wumpus Advisor

The methods used by a tutor to gather feedback are largely determined by the teaching environment and the subject matter that is being taught. The teaching environment restricts the options available to the tutor so that many techniques, such as jokes, are not available to a computerized tutor. Also, the subject matter lends itself to various feedback approaches. Geography is a field in which it is very difficult to evaluate the student's retention of the material without directly asking questions concerning that knowledge. (v)

The Wumpus game is a domain which is conducive to obtaining feedback by observing the student's behavior. This is true of most games environments as the game presents the student with problems to solve; an

---

(v) A geography tutor would be expected to ask questions such as, "*What is the capital of Brazil?*", which is a direct testing of the knowledge that is being taught. Just such a geography tutor has been developed by Carbonell and Collins in SCHOLAR.

Advisor need only observe the results. In such environments, it is important that the Advisor be careful to take full advantage of these observations and minimize the number of questions which it asks of the student. If the Advisor does resort to "testing" the student as a normal procedure, the student's enjoyment of the game is likely to be greatly diminished, and he is likely to resent the ministrations of the Advisor. For these reasons, the Wumpus Advisor relies almost solely on observations of the student for its evaluations of his knowledge of the game.(vi) The only way in which the Wumpus Advisor can be said to be eliciting feedback is the manner in which it defines (and perhaps alters) the game in order to create environments which are more conducive to learning and the evaluation of the student.

The above two spectrums are, of course, simplifications, but they help to emphasize two very important issues of tutoring strategies. Other issues that must be considered (and warrant further experimentation) are such issues as graphical versus verbal explanations and the types of instruction which are most effective (such as proof by contradiction, which is one the key methods of the Advisor for simple proofs). It is expected that Wusor III will be able to vary its tutoring strategies according to the student and will consider each of the above issues as well as other of such issues. These additions will increase Wusor's Learning Model to the dimensions described in Goldstein's proposal for the Coach Project (1977).

Conclusion

(vi) For the purpose of testing the validity of the Advisor's deductions, a useful experiment is to ask questions of the student to determine the correctness of the Advisor's estimation of his knowledge.

In designing Wusor II, there were several possible approaches to the problems of presenting the material and evaluating the student's response. The Advisor presents explanations of the results of the algorithms, instead of the algorithms themselves, as this seemed to be more conducive to teaching better ways to think about problems. Likewise, it was decided that the feedback from the student should be gathered primarily by observing the students behavior as this would not detract from the students enjoyment of the game and would help maintain a positive feeling on the part of the student toward the Advisor.

The English Generation Routines

Chapter 7

The English used by the Wumpus Advisor is generated in a conceptually clear manner. For each of the routines/rules of the Expert, there is a corresponding English routine/rule. If a routine of the Expert uses the results of a lower level routine, the corresponding English routine will likewise insert the text generated by the corresponding lower level English routine into its text. (1) In this fashion, it is a quite simple to write short functional English generation routines. The problem with such an arrangement is that the English thereby generated tends to be rather long and awkward. For this reason, various additions were made to the basic framework just described to improve on the output of the English routines. In the remainder of this chapter, we will describe those additions which were made to this simple approach.

A very important addition to the English routines was to make them aware of what the player already knows. In this fashion, the English routines can omit those portions of explanations which the player is thought to already know. This is not extremely important in the earlier phases of the curriculum, but it becomes more important later on. As the Wumpus Advisor teaches progressively more advanced rules, it relies increasingly on the simpler rules which have already been taught. If the English routines do not omit the lower level explanations, the explanations become quite long (and it becomes extremely unlikely that the student will read the explanation at all). The general solution to this problem is to have each English Routine first check to see if the player already knows all of the rules involved in its explanation. If so, then the English routine will not generate a complete explanation, but will instead only summarize its results. This simple check generally suffices to prune out the unnecessary parts of the explanations.

---

(1) The approach of having the explanation routines model the decision making routines is fully described in Stansfield (1975).

Figure 7.1

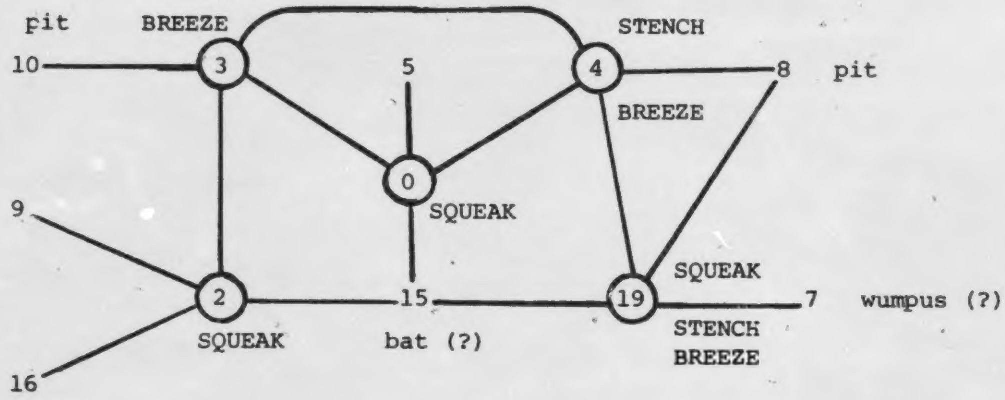

Circled caves have been visited.
Upper Case: Given Facts
Lower Case: Deductions

For example, a moderately advanced player was faced with warren shown in Figure 7.1 and he received the following explanation of why cave 15 was very likely to contain bats (P13).

> ***One of caves 15 and 7 contains bats. Likewise, one of caves 5 and 15 contains bats, and one of caves 16, 15, and 9 contains bats. As you have multiple evidence for bats in cave 15, it is very likely that cave 15 contains bats.***

This explanation presumed that the player already knew all the rules necessary to derive cave-sets, so that the explanation routines simply listed the cave-sets concerned rather than explaining each one. However, once the student model was modified to indicate no knowledge of the required rules, the following explanation was obtained:

> *One of caves 15 and 7 contains bats. Cave 19 must be next to bats because we heard squeaking in cave 19. Hence one of caves 15, 7, 8, and 4 contains bats, but we have safely visited cave 4. Also, cave 8 can not contain bats because if bats were in cave 8 we would have heard squeaking in cave 4 as cave 8 is next to cave 4. This means that one of caves 15 and 7 contains bats. Likewise, one of caves 5 and 15 contains bats. Cave 0 must be next to bats because we heard squeaking in cave 0. Therefore one of caves 5, 15, 4, and 3 contains bats. However we have safely visited caves 3 and 4. This means that one of caves 5 and 15 contains bats, and one of caves 16, 15, and 9 contains bats. Cave 2 must be next to bats because we heard squeaking in cave 2, and so one of caves 16, 15, 9, and 3 contains bats, but we have safely visited cave 3. This means that one of caves 16, 15, and 9 contains bats. As you have multiple evidence for bats in cave 15, it is very likely that cave 15 contains bats.*

The Wumpus Advisor would never have made such an explanation under normal circumstances (ii), but this example does point out the importance of pruning the English for even moderately advanced explanations. Without pruning, explanations quickly become so long that the the overall thread of the explanation is lost.

(ii) Normally the Curriculum would prevent the Advisor from explaining rule P13 before the student had thoroughly mastered all of the rules involved in explaining cave-sets, but for this example the Curriculum was disabled.

Another problem with the very simple model of the English routines is that such a model does not consider the idiosyncracies of the English language. For example, the English routines should not say "one cave away from" or "zero caves away from" (as was common in this paper), but rather it should say "next to" and "contains". There are a great many of such special cases (most of which center around the end conditions), and the English routines must have this knowledge programmed into them. Such problems can be handled with little deviation from the conceptual model if they are handled modularly.(iii)

In order to finely tune the English routines so that its output consistently sounds natural, each English routine should be written with an idea of the context in which it will be used. An example of this is the amount of introduction given by an English routine for a lower level English routine's explanation. If the lower level routine returns only a short summary, then only a short introduction (if any) is warranted. In contrast, if the lower level routine returns a lengthy response, a longer introduction is warranted, and, in some cases, the entire explanation should be restructured. Such issues can be resolved relatively easily by considering the curriculum, and the overall hierarchy of the rules. It is unrealistic to presume that the English routines which explain complex probability rules will ever be called before the student has mastered the simpler probability rules. With this in mind, it is reasonably safe for the more complex probability routines to presume that the student will have thoroughly mastered cave-sets by the time they are called. This

---

(iii) The solution to the problem stated above is simply a routine that converts "N caves away from" into its English equivalent.

reliance on the structure of the curriculum allows the context sensitivity to be static.

The addition of various checks to the conceptual framework originally described has changed what would have been very simple routines into relatively long and complex routines. An alternative method of improving on the English generated by simple routines is to use a standard English Generation Front End as shown in Figure 7.2. The Front End would accept the relatively "rough" output of the English Generator and then improve on it by performing such tasks as pronoun insertion, introduction modification, phrase replacement, etc.. This would have the effect of improving the overall quality of the English generated (as great effort has gone into developing such Front Ends) (iv) and would also make the task of generating English simpler. (Less knowledge of the English language would have to be programmed into the Advisor's routines as that knowledge is already programmed into the Front End).

Figure 7.2
Diagram of an English Generation System
Arrows Represent Data

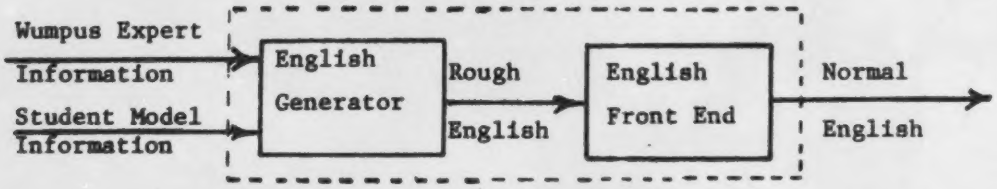

(iv) For more detailed information about English Generation Front Ends, see the work of McDonald, Novak, Simmons, Schank, etc.

Testing the Advisor

Chapter 8

The Wumpus Advisor is currently being tested to verify that is does, in fact, function as described. Also, a great many of the decisions which were made during the design of the Wumpus Advisor were based on the judgement of the author; these issues warrant further study to determine if the approaches taken were correct. Finally, it is important to determine whether the Advisor succeeds in teaching better ways to play Wumpus and, more importantly, better ways to think about problems. The tests can be divided into two general classes. They are local tests designed to test the ability of the Advisor to perform specific functions and global tests intended to test the overall ability of the Wumpus Advisor.

Local Tests

Local tests have been made with two different approaches to the problem. The first was a series of sessions in which students used the Wumpus Advisor in the presence of the author. In one such session, it was noted that negative logic (explaining the absence of warnings) was much more understandable than the approach then being used by the Advisor. Prior to this, the Wumpus Advisor would generate explanations like:

> *Cave 0 can not contain the Wumpus because we did not smell the Wumpus in cave 2. This means that cave 2 is more than two caves away from the Wumpus, cave 1 is not next to the Wumpus, and cave 0 can not contain the Wumpus.*

This (positive logic) seemed unnatural to students, and they responded much better to the negative logic which is currently used by the Wumpus Advisor. In the same situation, the Wumpus Advisor would explain:

> *Cave 0 can not contain the Wumpus because, if the Wumpus were there, we would have smelled the Wumpus in cave 2.*

There were approximately ten such sessions of roughly three hours duration, and, based on subjective evaluations of the sessions, it was determined that in general the Advisor functioned as desired, but it was not possible to determine how well the Wumpus Advisor would function in the wide range of situations it would eventually face.

The next group of experiments to test the Wumpus Advisor's ability to function as planned were conducted with a Synthetic Student.(i) The Synthetic Student (Syndi) was designed to play the game of Wumpus in conjunction with the Wumpus Advisor (Wusor). Syndi is a perturbation of the Expert of Wusor with the main change being the addition of conditional statements which allow Syndi to play with some subset of the full Expert. Syndi has a Knowledge Model of her own and her Expert will not apply any rule unless the Knowledge Model indicates that she possesses knowledge of said rule at that time. Using Syndi, a great deal of debugging of Wusor was achieved with relative ease. Wusor's analysis of the student's move involves several inherent assumptions and Syndi was designed to test Wusor's analysis given that these assumptions are true. The assumptions are:

1. The student does not apply paradigms unknown to the Wumpus Advisor.
2. When faced with several equally desirable choices for moves, the player will choose that move which is closest (and hence easiest to get to).

(i) The idea of using a Synthetic Student to test an Advisor program was originally described in Goldstein's Coach Proposal (1977).

3. The player is consistent in his application of rules. He does not have partial or incomplete understanding of a given rule (the undesirability of this assumption is decreased by the almost atomic rules of the Expert).

4. The player does not make mistakes. Wusor's relative caution is necessary because this assumption is obviously false.

5. The player comprehends the explanations given by Wusor.

Syndi was designed to satisfy all of the above assumptions. Also, as she had a distinctly diagnosable knowledge of the game at all times, it was quite simple to verify any results obtained by the Advisor. Several experiments were run with Syndi and Wusor, and they can be grouped into three types;

1. Experiments in which Wusor initially presumed Syndi had no knowledge of the game and Syndi played with a constant knowledge state. Ideally, Wusor should advance his estimation of Syndi's knowledge to exactly correspond to her actual state and not beyond.

2. Experiments in which Wusor initially presumed Syndi had a complete knowledge of the game and Syndi played with a constant knowledge state. Ideally, Wusor should degrade his estimation of Syndi's knowledge to agree with the actual state.

3. Experiments in which Syndi started as a new player with no knowledge of the game, and "learned" rules as Wusor explained them to her. Ideally, Syndi's knowledge should increase to complete knowledge of the game with Wusor's estimation lagging a little behind each change.

These experiments are relatively complete and thoroughly test the Advisor's functioning given that the above assumptions are true. In experiments of Type 1 (advancing to a fixed knowledge state), it is acceptable to always start with Wusor presuming Syndi has no knowledge of the game because Wusor must advance his estimation through all of the phases up to the desired phase to achieve the desired result. This also justifies starting all Type 2 experiments (degrading to a fixed knowledge state) from an initial estimation of complete mastery of the game. Experiments of the first two types serve to test the existence of situations which will allow Wusor to make deductions as well as testing the correctness of said deductions. The third type of experiment is necessary to insure that there are situations which allow Wusor to tutor a student from the very basics to the most advanced skills in a moderate amount of time.

When conducting the above experiments, several modifications were found necessary. Many of the experiments with Syndi did not achieve the desired results on their first run, which required that Wusor be modified in light of the preliminary results. Then the process would be repeated. In this section, we will discuss the experiments which were run, the problems which arose, the solutions which were found (if any), and the final results achieved. All this will be in general terms with the specific results achieved shown in Appendix F.

The first experiments were run with Syndi having a complete knowledge of the game and Wusor initially presuming that she had no knowledge of the game. It was found that Wusor advanced the student rapidly with respect to the Wumpus domain, but unreasonably slowly with respect to the bat/pit domain. Because of this, several modifications were required before Wusor would correct its estimation of the student's knowledge to complete mastery of the game within a reasonable amount of time. The first problem was that Syndi would risk bats far more regularly than pits (as bats are not necessarily fatal) so that Wusor had very few chances to observe the application of the probability rules with respect to pits. This required the modification of the student modelling components so that the bat/pit *Familiarity Values* would be associated together once Wusor had determined that the player had mastered the rules to determine which caves are safe (i.e. Wusor would presume a transference of bat/pit knowledge once the player had mastered Phase 1). The next modification was required because there were often long periods in which there were no situations which would allow Wusor to deduce a mastery of probability rule P11 for bats/pits. In order to overcome this difficulty, Wusor was given the ability to modify the initial game situation to create ideal situations appropriate for a player of the given phase. These ideal situations are described in Appendix C and illustrated in Appendix D. The desired results were achieved by combining the above modification with another modification; the criteria for deducing that the player had mastered a given skill was lowered somewhat. Previously, the situation had to have been such that the selected move was better than some other move with respect to a single danger and identical (no differences) with respect to the other two dangers. Then Wusor would note that the player had given the appearance of mastery of those rules which were required to justify the superiority of the better cave for the relevant danger. This was relaxed so that, if there were a single deciding danger which made the selected move superior, the two moves had to be one of the following with respect to each of the two remaining dangers: 1) identical, 2) such that the player should have perceived that they involved the same risk, or 3) such that the player should have perceived that the better move involved a greater risk. These two modifications, along with fixing various outright bugs, allowed Wusor to advance the student to complete mastery of the game in roughly thirty games (a figure which is considered acceptable for a tutor designed to work with children).

The next series of experiments were such that Syndi used a fixed subset of the expert's knowledge and Wusor started with an initial estimation of no knowledge of the game. In these experiments, Wusor advanced rapidly to the actual state (as it did before), but rather than stopping, it continued to advance the student slowly, albeit surely. This required the addition of the more complete analysis for probabilities described in Appendix B. Of course the original experiments had to be run over again, but it was found that in all cases Wusor would advance to the correct knowledge state and then make no further presumptions of the student's knowledge. The advances proceeded at an acceptable rate. The only non-optimal performance by Wusor was due to problems in implementing Syndi. In developing Syndi it was found that it was necessary that she have some minimal knowledge of the game at all times in order to model even beginning players. The required knowledge was an understanding of what it meant to have encountered a danger, what it meant to have safely visited a cave, and a rudimentary understanding of what it meant to shoot into a cave. The problem was that the rudimentary understanding of the shooting principle (P15) would occasionally cause Wusor to advance the student slightly beyond the actual knowledge state, assuming a full mastery of rule P15 and the rules associated with it when, in fact, there was only a partial mastery (i.e. Assumption 3 was violated).

The next set of experiments were of the Type 2 described earlier. Syndi was started with a fixed, but limited knowledge of the game, while

Wusor initially assumed that she had complete mastery of the game. The ideal was that Wusor would degrade his estimation of Syndi's knowledge until it corresponded to the actual knowledge and then make no further change. Several experiments were run with different knowledge states for Syndi. Many of the same problems that were encountered in advancing the student were encountered in degrading the knowledge estimation. It was found that there were not sufficient game situations which would clearly reflect on the application of rules P11 for bats/pits (as in advancing the student for the same rules). The problem was that the ideal situations which had been introduced to aid in the advancement of the student were always one level more advanced than would be necessary to degrade the student's knowledge model. Fortunately the implementation of the ideal situations had been such that Wusor was only able to create an ideal situation approximately one fourth of the time (ii), and this allowed Wusor to try to create an ideal situation for a student of the next lower level (and the next more advanced level) if it failed to create the ideal situation for the given level. This was sufficient to overcome this problem.

Another problem that was encountered when running the Type 2 experiments was that Syndi sometimes would not risk pits often enough to allow Wusor to analyze her knowledge of the pit rules. This problem had been encountered earlier when testing Wusor's ability to advance the student and had been solved by assuming that knowledge in one of the bat/pit domains would be transferred to the other domain. A solution to

---

(ii) This was done to restrict the possibility of the student deducing that Wusor was modifying the game.

the problem would have been to completely tie the two domains together with a resulting lack of flexibility (since Wusor would not be able to diagnose failures to transfer knowledge from one domain to the other). It was decided that the decrease in flexibility was not warranted in light of the fact that decreasing Wusor's estimation of the student's knowledge was not the normal mode in which Wusor would function. This same trade off came up when it was noted that Wusor would sometimes (though rarely) note that the player had demonstrated a knowledge of rules which the student did not, in fact, possess. This was caused by the reduced requirements for demonstrating knowledge of a rule which were implemented while running the Type 1 experiments. In particular, because Wusor assumed that the student would recognize that a move was worse because of a certain danger (and in this case that was a false assumption), and therefore conclude that he he must have seen the virtues of the move with respect to another danger (since he chose the move). This was determined to be a minor problem because the problem was very unlikely in normal situations (since Wusor would not be that far off in his estimation of the student's knowledge) and because Wusor eventually corrected for the error. To return to the more stringent requirements would decrease Wusor's ability to advance the student (his normal mode) without significantly affecting Wusor's ability to decrease his estimation of the student's knowledge.

The above two sets of experiments showed that Wusor was completely accurate in its ability to increase its estimation of the student's knowledge to the appropriate level (and not beyond) when its intrinsic assumptions were valid. It was also found that, when the assumptions were violated (in particular, the assumption that the student would not have partial or incomplete understanding of a rule), Wusor could diagnose the general level of the student, though its estimation of the student's knowledge of particular rules was not always reliable. It was also found that Wusor had a limited ability to decrease his estimation of the student's knowledge, though this mode was not nearly as reliable. In light of these results, Wusor's initializations of the student's knowledge were always kept on the conservative side.(iii) Having verified that Wusor's analysis of the student's actions were working, it was then appropriate to insure that Wusor could give advice on all the different aspects of Wumpus play. This was to insure that, if the student could understand Wusor's explanations, it would be possible for him to become an expert in Wumpus play without any outside help (i.e. there were not any holes in Wusor's curriculum).

The final set of experiments were of Type 3, in which Syndi started with a minimal knowledge of the game, but would learn each rule as soon as it was explained to her. After a certain level, Syndi would begin transferring knowledge between the bat and pit domains. Wusor started with a correct estimation of Syndi's ability. Of course, Syndi's ability to learn is rather artificial. When Wusor notes that he has just explained a rule to the student, Syndi will be so informed as well (thereby allowing her to begin applying the rule in question). Another technical problem was to decide how Syndi should modify her behavior when

---

(iii) The claim that Wusor's initializations of the Student Model tends to be conservative is justified by the fact that, in real usage, Wusor typically starts the student at a certain level and quickly advances the student to a more appropriate level; after which advances tend to come more slowly.

she first learned a rule. Should she recalculate her entire data base or change it as her attention was called to it. This was very important as human players do not take either extreme (but, instead, do a little bit of both). The general tack that was taken was to emulate a human player as much as possible. For example, if a player did not previously understand the significance of warnings/nowarnings, it is unlikely that he would have kept a record of this information. When he did learn their meaning, he would have to revisit a cave before he would able to use the information of its warnings. So it was with the data base with Syndi. This is not true of the more complicated rules as they simply involve better usage of the available information. This aspect of human players was also modelled in Syndi.

In the Type 3 experiments it was found that Wusor was able to find situations to advise the student for all levels of knowledge of the bat/pit domain (again, this was largely due to Wusor's ability to modify the game to create ideal situations for bat/pit knowledge), but Wusor was not able to ever advise the student with respect to the rules of Phase 2 of the Wumpus. This was because of a bug which made Wusor unnecessarily restrictive in the situations where it would advise the player about these rules.(iv) This bug was fixed, and Wusor then began explaining to Syndi the rules of Phase 2 for the Wumpus. Once this obstacle was overcome, it was found that the situations in which Wusor would explain the rules of Phase 3 for the Wumpus were extremely rare. This problem was fixed by creating a new Combination Rule, Combination Rule Seven. The

---

(iv) The bug was mutually exclusive requirements prior to the explanation of Combination Rule 5.

above two modifications allowed Wusor to advance Syndi from the most basic levels to the most advanced levels by first explaining a rule and then observing when Syndi applied the rule. In Appendix G there is the first part of the protocol taken in a Type 3 experiment. It is indicative of what all the protocols look like.

The last three of the five assumptions are most questionable. The above experiments indicate that Wusor functions as desired even when the third assumption (the assumption about incomplete understanding of rules) is violated. The fourth assumption that the player does not make mistakes is obviously invalid, but, because of this, Wusor was designed to be cautious in its conclusions (thereby decreasing the detrimental effects of mistakes). The remaining questionable assumption is the assumption that the player understands Wusor's explanations. There is evidence that Wusor is, in fact, effective at tutoring students in his game strategies. While the above experiments were being carried out, Wusor was accessible to the general community at the Artificial Intelligence Lab so that anyone who was interested could play with Wusor. Roughly 15 students played with Wusor extensively with reasonable results (v), but there was one disturbing anomaly. It was known that there were more situations which allowed the student to demonstrate knowledge of the Wumpus domain because Wusor advanced his estimation of Syndi in this domain much more quickly than in the bat/pit domain. However, when Wusor was being used by actual students it was observed that knowledge of the bat/pit domain far exceeded the Wumpus domain beyond Phase 1. It was feared that this indicated that Wusor assumptions were not valid, but then it was discerned that Wusor was unable to tutor students on the more advanced Wumpus rules. This, in itself, is indicative that students do tend to adopt those strategies explained by Wusor, though stronger evidence will be gained if students advance more quickly in the Wumpus domain now that Wusor can advise them in these skills.(vi)

(v) This is the subjective evaluation of the author based on his study of the complete protocols kept by Wusor.

When a large tutorial program, such as Wusor, is designed with such varied behavior, it is extremely difficult to determine whether or not the program functions as desired. However, this testing can be vastly simplified through the use of a Synthetic Student. Syndi served to test Wusor in closed experiments where all variables could be controlled (Syndi never had a bad day). Of course, experiments with Synthetic Students do not conclusively show the effectiveness of a tutor, but they do serve as a wonderful debugging aid. With a Synthetic Student, it is possible to determine whether or not the program functions as desired; the determination if the correct approach was taken has to be determined through other experiments which are of a more global nature. We will go on to describe these global experiments in more detail.

Global Testing

The Global Testing is the next step in testing the Wumpus Advisor. One of the goals of this phase of experimentation is to test out the relative merits of different tutoring strategies, such as graphical versus verbal. However, the main goal of these experiments is to

---

(vi) The reader can also verify that Wusor's explanations are comprehensible by studying the scenarios in Appendix E and Appendix G.

determine how well the Wumpus Advisor actually teaches better game playing (if it does, in fact, help students master the game). The experiment would probably be a rather simple arrangement involving a total of roughly ten students. Five of these students would be allowed to play the game of Wumpus (as modified in this paper) without the Advisor. The remaining five would be allowed to play the game with the advice of the Wumpus Advisor. At the end of the testing period (which would almost certainly extend over a period of about two weeks to a month), the students would be tested on their understanding of the game. There would be the objective results of the percentage of games won and lost when the students were allowed to play without any help, but there should also be subjective results. The students should be tested on their ability to make deductions, their ability to estimate dangers, and their general understanding of the game. It is expected that the students who played with the Advisor will not only play better (as indicated by their percentages of wins and losses) but also have a better understanding of the game.

The final experiment is to test the Wumpus Advisor's ability to teach better ways to think about problems. Only a general outline of the experiment will be given as the details will be dependent on the results of the previous tests. The experiment should involve about 150 students to allow for statistical evaluation. At the start, all 150 students will be given a test that measures their ability to think about problems (probably an I.Q. test with emphasis on thinking ability). After the test, the students will be separated into homogeneous groups of 50. The first group of 50 students will be left alone for six months. The second group will play the game of Wumpus without the Wumpus Advisor. The last group will play the game with the Wumpus Advisor. The experiment will continue for six months with the active students playing Wumpus for a specified period each week. Finally, all the students will be given a variation of the original test. In this manner, it should be possible to test both the game of Wumpus and the Advisor in their ability to teach better thinking habits.

The large size of the groups to be tested is essential as it is doubtful that any one force can be that significant in improving thinking habits. With a smaller sampling size it is possible that the small increments in thinking ability would be lost in the normal noise that accompanies such tests. It is also possible that the results would not be conclusive even with such a large sampling, but there is no way to predict this. A significant number of the subjects in the standard curriculum of the public schools are justified by the hope that mastery of the given field will help the student develop better ways to think about problems. However, no one seems to have made a significant effort to verify these claims of improved thinking habits, and so it is very difficult to estimate the sample size which will be necessary to arrive at conclusive results. The proposed test should help to clarify how large a sample is required to arrive at conclusive results in such a test.

Conclusion

The Wumpus Advisor has been thoroughly tested to debug it and fine tune its tutorial ability. Concurrently, consideration was given to the various programming decisions which were made according to the judgement of the programmer due to the lack of concrete evidence. The Advisor has been determined to be "ready," testing is in progress to determine the Advisor's ability to teach better game playing. This experiment will be on a relatively small scale because it is expected that the results will be conclusive in any case. However, when the Wumpus Advisor is tested on its ability to actually teach better ways to think about problems, the results are not expected to be as obvious. For this reason, a larger sample size is necessary if the experiment hopes to be conclusive. Finally, based on the results of these experiments, further work leading to Wusor III could be warranted.

The Conclusion

Chapter 9

A new generation of Computer Aided Instruction (CAI) is evolving which transcends past limitations by using Artificial Intelligence (AI) techniques. These new programs are AICAI programs since they incorporate the concepts of Artificial Intelligence while maintaining the goals and objectives of Computer Aided Instruction. AICAI is a field which contains many interesting and challenging problems, the Wumpus Advisor being just one example.

The Wumpus Advisor applies many of the techniques developed in Artificial Intelligence. It centers around a rule-based Expert which uses heuristics to arrive at results that are both reliable and understandable to humans. The rules of the Expert are organized into a curriculum which is taught by the Wumpus Advisor as appropriate. The rules of the Expert are also the foundations for the Student Knowledge Model. The Student Learning Model interprets the Student Knowledge Model and determines which rules the student is presumed to have learned and which rules are acceptable for teaching. In making these determinations, the Student Learning Model makes extensive use of the Student Learning Vector. The Expert's rules are also the basic framework around which the English generation routines are built. The fact that both the Student Model and the English Generator are designed around the rules of the Expert allow the English routines to prune their English according to the student's knowledge. The logical grouping of the Expert's rules in the curriculum also allows the English routines to be context sensitive without lengthy computations. The essence of the improved Wumpus Advisor is the different rules of the Expert and their interrelationships.

The main limitations of the original Wumpus Advisor were due to the lack of a clear understanding of the knowledge to be taught. Once this knowledge had been broken down into specific items, the development of the remainder of the Wumpus Advisor was relatively straight forward. In more interesting problem domains (such as chess), it will probably not be as difficult to develop the rules for the expert, because the problem domain will have been analyzed before the program is begun. The main task will be to break down the knowledge of the problem domain into rules. This will allow more effort to be directed into those areas which are of more theoretical interest (such as the Student Model).

The next logical step beyond the Wumpus Advisor will be a program that teaches skills in a problem domain that can not be analyzed with absolute certainty. An important asset of the game of Wumpus was that it was possible to discern a single best move and explain why. In problem domains such as chess, it is not possible to know absolutely the quality of most moves. This would handicap an advisor since there would always be the possibility that the expert would not give correct advice. However, this is a handicap shared with human teachers. An advantage of more complex problem domains is that the Advisor will be able to teach principles and concepts which can not be represented within the limited domain of Wumpus. For example, the formal methods of decision theory could not be taught effectively within the Wumpus domain, but it could be taught in more complicated domains. An advisor which functioned in a more complicated domain could be built along the general lines of the

Wumpus Advisor, but its Student Model would require significantly more effort. An advisor for a more complicated domain would need to be more flexible in its explanations to the student since the advisor would not be as certain of the results of its expert. The primary source of this increased flexibility would be the Student Model.

Appendices

Modular Diagram of the Wumpus Advisor

Appendix A

Each sub-module of the Wumpus Advisor has a unique prefix which identifies all the routines of that sub-module. These prefixes are used to identify the sub-modules in the flow chart. Each sub-module performs a particular task which is described in the remainder of this Appendix.

The Wumpus Executive

WE- is the prefix for the Wumpus Executive routines.

WG- is for the Wumpus game.

The Wumpus Advisor

WA- is the prefix for the Wumpus Advisor routines.

WAW- is for the routines of the Wumpus Advisor which keep track of the route taken by the player.

WAD- is for the routines which compute the distance to different caves from a given starting point. They also compute routes, etc..

WAM- is for the Wumpus Advisor routines which modify the game and create the various starting situations.

The Expert

XX- is for the executive for the Wumpus Expert. It activates the other sub-modules and then combines their results.

XD- is the prefix for the routines which maintain the data base for the Expert. They also compute the "MORE-THAN" and "EXACTLY" values.

XS- routines keep track of the different cave-sets, noting which are complete, etc..

XR- routines identify those caves which must be safe based on global consideration of the cave-sets for a particular danger and the number of the dangers in the warren.

XP- is the prefix for the routines which compute the probabilities for the different dangers.

The Move Comparer

CX- is the prefix for the Result Explainer routines. They return the rules which are necessary to justify the results of the Expert.

CM- is for the Explanation Comparer, which compares the rules returned by the Result Explainer and determines how familiar the student is with the rules in question.

The Psychologist

PS- is the prefix of the Psychologist.

The Student Model

SK- is the prefix of the Student Knowledge Model. These routines keep track of the student's knowledge of the different rules.

SL- is the prefix of the Student Learning Model. These routines estimate how quickly the student learns material, forgets same, etc..

SC- is the Student Model Critic. It adjusts the Student Learning Model as appropriate.

SF- routines maintain the disc files on the different users.

The English Generation Module

EX- is the prefix for the routines which generate English to explain the results of the Expert.

EG- routines generate clauses in English that convey a particular concept (such as "N caves away from").

Utility Routines

GP- is the prefix for general purpose utility functions which expand the abilities of LISP.

G- routines are the I/O routines of the Wumpus Advisor.

GQ- routines query the student with standard questions and return his response once it is in an acceptable form.

GM- are general purpose routines which do mapping (such as a routine to do MAPC non-destructively).

GC- routines do the processing of the circular list which are used in English generation and FIFO stacks.

PLAYER is an extraneous variable in an otherwise conceptually clear framework. This anachronism will be eliminated as soon as possible.

87

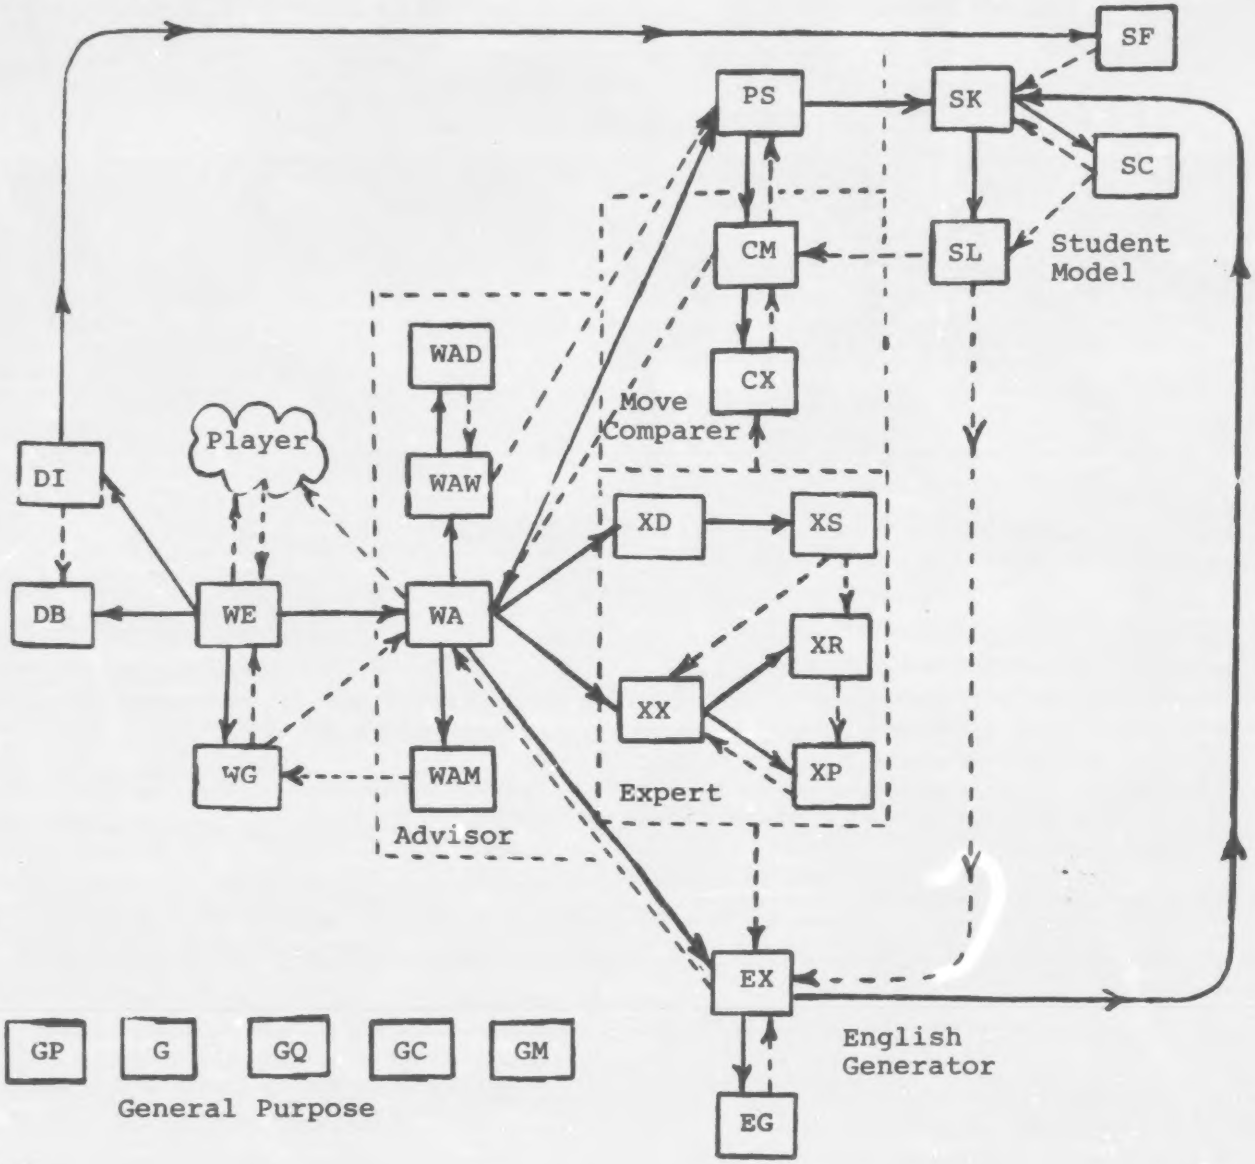
SF
PS
SK
SC
CM
SL
Student Model
WAD
CX
Move Comparer
Player
DI
WAW
XD
XS
DB
WE
WA
XR
XX
WG
WAM
XP
Expert
Advisor
EX
English Generator
EG
GP
G
GQ
GC
GM
General Purpose

Modular Diagram of the Wumpus Advisor
Dotted lines represent significant flows of information.
Solid lines illustrate the hierarchy of control.

Move Comparisons

Appendix B

The Move Comparer is responsible for taking two moves and justifying their relative merits with respect to a single danger with specific rule numbers. For this purpose there are two separate sub-modules to perform this task. These two sub-modules are the Result Explainer and the Explanation Comparer. First, the Result Explainer determines which rules were used by the Expert to arrive at a specific result.(i) Then the Explanation Comparer takes two results and compares the different rules involved. In this fashion, the Move Comparer can return the rules which justify the difference between two possible moves with respect to a single danger. The Move Comparer also performs various other functions; primarily, the Explanation Comparer consults the Student Learning Model and determines if the student knows the rules in question, are they acceptable for teaching, etc.. This is a cursory overview of the process, but we have overlooked many of the problems of such comparisons. In the remainder of this Appendix we will look at these issues more closely.

When the Wumpus Advisor is comparing two different moves it must occasionally consider which of the probability rules is involved. This is true if the Wumpus Advisor is analyzing the student's move to update his Student Model, or if the Wumpus Advisor is determining if it can explain a particular good move to the player. In any case, the Move

---

(i) In this light, the Result Explainer is structured almost exactly like the Expert.

Comparer must be able to decide which probability rules are applicable in different situations.

The Wumpus Expert only considers the most accurate probability rules in its analysis, but the Wumpus Advisor must be able to determine when less advanced probability rules are involved in a particular comparison. For example, in Figure B.1 there is a common situation that is presented to moderately advanced players. Cave 8 is very likely to contain a pit because of the double evidence of a pit (i.e. the *breezes* which were felt in caves 3 and 12). This makes it rather unlikely that either of caves 4 and 18 contain a pit since a pit in cave 8 would explain the *breezes* in caves 3 and 12. The rule which justifies the high probability of a pit in cave 8 is P13. The rule which justifies the low probability for caves 4 and 18 is rule P14. Now, if the Wumpus Advisor were comparing the two probabilities for caves 4 and 8, it might note that rules P13 and P14 were both involved in the comparison. However, a student could know only rule P13 and not rule P14, and he would still choose the better move. Correspondingly, the Wumpus Advisor could explain why cave 8 was a bad move by simply explaining how P13 applied.

Figure B.1

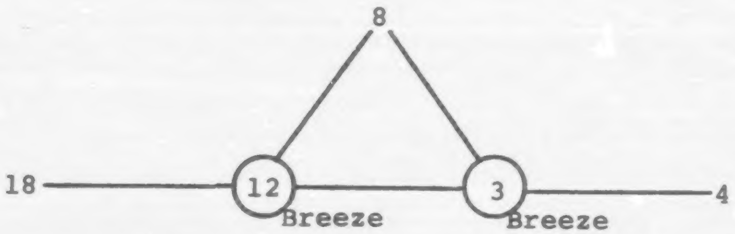

Such situations require that the Wumpus Advisor (in particular, the Move Comparer) make a detailed analysis when two probabilities are being compared. For this purpose, the Wumpus Expert computes the probability for every cave with respect to each probability rule. It then chooses the prefered probability (the most reliable) and saves all of the different probabilities. The first step of the Wumpus Advisor's analysis is to insures that one cave actually does have a lower risk than the other move.(ii) It then determines which of the following cases applies by checking them in order.

A) If the better cave (*Better*) is safe, then the only applicable rules are the rules which make *Better* safe. This is because the presumption is that caves are moderately dangerous unless there is some evidence indicating that they are safe.

B) Correspondingly, if *Worse* is certain to contain the danger, then the only applicable rules are the rules which show that *Worse* contains the danger (normally L0). This is because the presumption is that *Better* is moderately unsafe.

C) If both of the above two rules are applicable, then choose the simpler of the two explanations (least advanced) as the necessary explanation.

D) If the comparison of the two P11 probabilities will justify the quality of the *Better* move, then P11 is the only applicable rule.

E) If P12 is applicable for the *Better* cave's probability, then P12 is the only rule that is applicable.

F) If the P13 probability for *Worse* is greater than the P11 probability for *Better*, then P11 and P13 are the appropriate probability rules.

G) If the P14 probability for *Better* is less than the P11 probability for *Worse*, then the probability rules which are applicable are rules P11 and P14.

H) If none of the above apply, then use the simplest probability

---

(ii) This comparison is based on the "preferred" probabilities of the Wumpus Expert.

rules which will justify the "preferred" probability.(iii)

In this fashion, the Wumpus Advisor determines the simplest possible rules which can justify the difference between two caves. This information is used in analyzing the student's move as well as in determining whether or not the better move is suitable for explaining for the student. Once the Wumpus Advisor has discerned that a move can be explained to the student, it then determines the best rationale to explain to the student. In the example above, the Wumpus Advisor might decide to explain to the student that cave 4 is a better move. Depending on Wusor's estimation of the student's knowledge, Wusor would choose to explain either rule P13 or rule P14. Wusor's explanation would center on rule P14 if the student already knew rule P13, and the explanation would center on rule P14 if the player was not yet advanced enough to allow the explanation of rule P14.

Once Wusor has decided to explain the comparison between two moves and has selected rationales to be explained, it must then create an explanation with the appropriate emphasis. The points of emphasis of such explanations are shown in Table B.2.

Table B.2

Rationale For Better Cave	Rationale For Worse Cave	Emphasis of Explanation
Rule P11	Rule P11	The Cave-Sets involved
Rule P12	Not Rule P12	The explanation of P12
Rule P14	Not Rule P14	The explanation of P14
Not Rule P13	Rule P13	The explanation of P13
Rule P15	Not Rule P15	The explanation of P15
Any Other	Any Other	Both Probabilities

(iii) If a cave is only a member of one cave-set, then the probabilities for P11, P13, and P14 will be equal in most cases. Therefore, P11 is applicable for that particular cave.

Wusor will generate the explanation which first fits. Examples of some of these explanations are in Appendix D.

Analysis of Game Modifications

Appendix C

Before the start of a game, the Wumpus Advisor will attempt to modify the game to create a situation which is thought to be a challenging problem to the student. This is done by looking for a starting cave within the warren which meets all the requirements of the situation desired. Typically, the requirements will involve the warnings at the cave in question as well as the warnings in one of its neighboring caves. If a cave is found that meets all of the requirements, the Wumpus Advisor will move the starting cave to the desired cave. Then, the player will be informed about his starting location (the one selected by the Wumpus Advisor) and what the neighboring caves are that he can move to. Once the player has selected a cave (at random since he does not have enough information to choose wisely), Wusor will transpose the caves of the warren so that he moves to the neighboring cave selected by the Wumpus Advisor. The requirements for starting situations vary for students of different phases. These requirements are shown in Table C.1.

Table C.1

Phase	Starting Cave Requirements	Neighbor's Warning	Requirement On Overall Situation	Teaching Point
0	No Warnings	Some Warning	None	Backtracking
1	Bat or Pit Warning	Other Warning	Neighbor in Common	L3 and L5
2	Bat or Pit Warning	Same Warning	Smaller Cave-Set	P11
3	Bat or Pit Warning	Same Warning	Neighbor in Common	P13
4	No Wumpus Warning	Any Warning	None	P13 (Wumpus)

Examples of each of these situations are shown in Appendix D. There are not acceptable starting caves in many warrens, so that the Wumpus Advisor only creates these situations roughly half of the time. This, when added to the fact there are various possible starting configurations for each phase, should be sufficient to prevent the student from ever recognizing the similarities of the starting situations before he is in a new phase. (iv) However, even when the Wumpus Advisor does not modify the game to create an ideal situation, it will transpose the player's first move if it is to a cave containing a danger, thereby preventing the player from ever losing on the first move (which can have very bad effect on morale). This ability to transpose moves can easily be extended to give positive or negative reinforcement as necessary.

(iv) This is a hypothesis which has not yet been verified through experimentation.

Examples of Ideal Situations and Probability Compa

Appendix D

This appendix contains examples of the "ideal situations" that were created for students in different phases along with the corresponding explanations. The requirements for the situations are given in Appendix C and the structure of the explanations is given in Appendix B. These are actual situations created by the Wumpus Advisor and the actual explanations given.

A beginning player (Phase 0) who is being taught the rules would be faced with a situation very similar to this.

Figure D.1

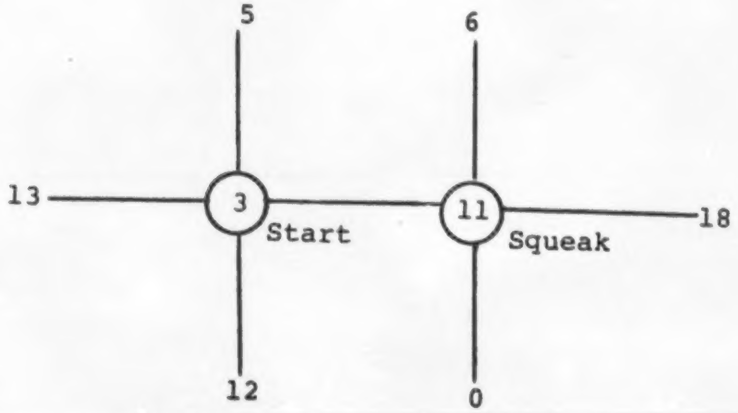

This situation would give rise to the following explanation if the player asked to move to cave 18.

> Brian, did you know that we can backtrack to caves that we have already visited? Yes
>
> Oh,... Well then why not go back to cave 3 where we didn't get any warnings at all. From there we can safely go to cave 13 without risking the bats. It is not wise to visit caves with bats because while THEY will not harm us they will carry us to another cave which could contain a pit or the

Wumpus.

Would you like to go to cave 18 anyway?

A player who was slightly more advanced (Phase 1) and presumed to have gotten a basic understanding of the rules, would be faced with the following situation.

Figure D.2

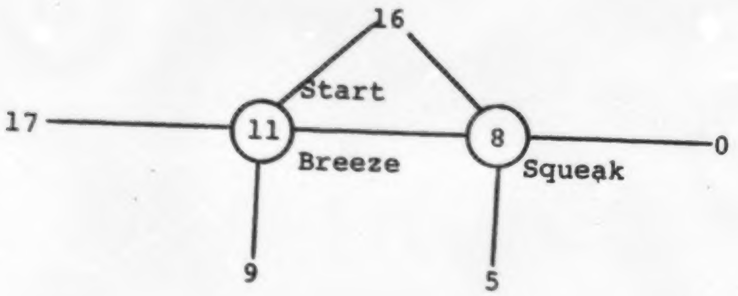

This situation would give rise to the following explanation if the player asked to move to cave 8.

> Brian, we don't need to risk bats in cave 8 as cave 16 is safe. Cave 16 can not contain bats because, if bats were there, we would have heard squeaking in cave 11. It is best to avoid bats as they could carry us to a fatal cave.

Would you like to go cave 8 anyway?

A player who is presumed to have mastered the rules necessary to discern which caves are safe (Phase 2), would be faced with a situation similiar to this.

Figure D.3

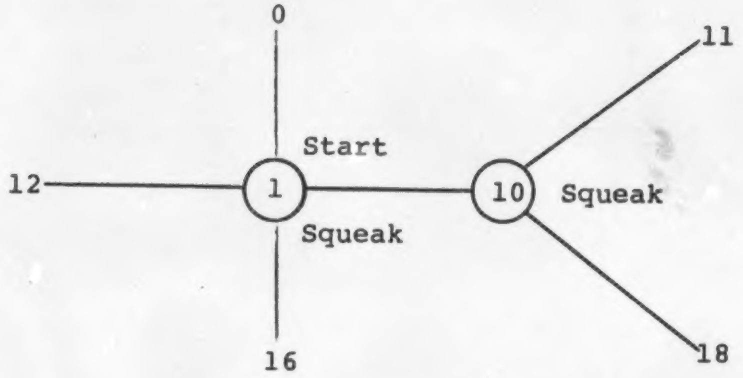

This situation would give rise to the following explanation if the player asked to move to cave 18.

> Brian, it isn't necessary to take such large risks with bats. It is true that one of caves 0, 12, and 16 contains bats because we heard squeaking in cave 1, but one of caves 18 and 11 contains bats because we heard squeaking in cave 10. This makes it less likely that cave 0 contains bats. Therefore, Brian, we might want to explore cave 0 instead.
>
> Would you like to go cave 18 anyway?

A player who has mastered the simple probability rules would be given a situation like this.

Figure D.4

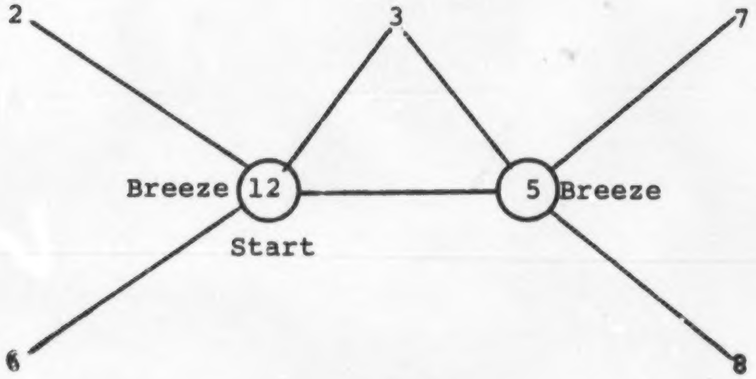

This situation would give rise to the following explanation if the player asked to move to cave 3.

> Brian, it isn't necessary to take such large risks with pits. One of caves 3, 7, and 8 contains a pit. Likewise, one of caves 3, 2, and 6 contains a pit. This is multiple evidence of a pit in cave 3 which makes it probable that cave 3 contains a pit. It is less likely that cave 8 contains a pit. Hence, Brian, we might want to explore cave 8 instead.
>
> Would you like to go to cave 3 anyway?

In the course of playing a game, a player of this level could be faced with a situation like this.

Figure D.5

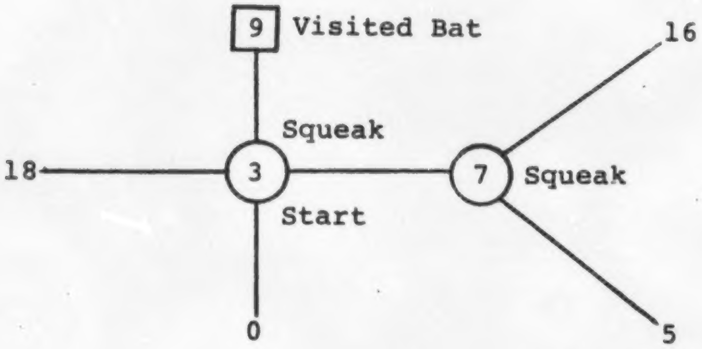

This situation would give rise to the following explanation if the player asked to move to cave 16.

> Brian, it isn't necessary to take such large risks with bats. It is true that one of caves 18, 0, and 9 contains bats, but cave 9 contains bats because we were picked up by bats when we visited it before. This explains all the evidence for bats in caves 18, 0, and 9. As we do not have any evidence of cave 0 containing bats, we can presume that it is unlikely that cave 0 contains bats. There is evidence of bats in cave 16 which makes it a more dangerous cave. Therefore, Brian, we might want to explore cave 0 instead.

This is an explanation of rule P12 for bats as described in Appendix B.

Sample Wumpus Advisor Scenario

Appendix E

In the following scenarios, both Wusor's statements and the player's responses will be indented to the actual fields used; the player's response will be underlined. These scenarios have been constructed by combining advice actually given by Wusor, but they are not realistic because Wusor would never give advice so often and because there is not a consistent Student Model throughout.(v) This Appendix is intended to give the reader some idea of Wusor's ability to give advice in a wide range of situations. In this regard, it is not all inclusive since Wusor has given advice in many situations far removed from the ones shown.

> Hello, my name is Wusor VII. What is your name? JIM STANSFIELD
>
> Do I have your name right, Jim Stansfield? Yes
>
> Would you like instructions? No
>
> Just a second while I erase my blackboard.
>
> In this game, Jim, there will be 20 caves all total, three pits, and three caves with bats.

In fact a new player would not be put in a warren with so many bats and pits.

> O.K., now I will draw up a new warren.
>
> We are now at cave 6.
> The neighboring caves are caves 3, 12, and 5.
>
> What now? 12

(v) Wusor's behavior is highly dependent on the current Student Model, and the following scenarios have been designed with a successively more advanced Student Model. In fact, the Student Model would not be advancing in such a manner because of the consistent mistakes by the player.

We are now at cave 12.
The neighboring caves are caves 1, 8, 18, and 6.
Squeak. I hear bats they must be in one of the neighboring caves.

What now? 1

Figure E.1

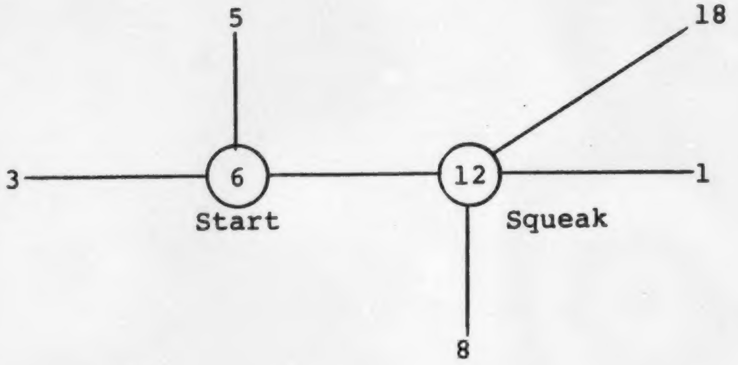

Jim, did you know that we can backtrack to caves that we have already visited? No

Well we can, so why not go back to cave 6 where we didn't get any warnings at all. From there we can safely go to cave 3 without risking bats. We should avoid bats because they could drop us in a cave with a pit or the Wumpus.

Would you like to go cave 1 anyway? No

We are now at cave 12.
The neighboring caves are caves 1, 8, 18, and 6.
Squeak. I hear bats.

What now? 6

We are now at cave 6.
The neighboring caves are caves 3, 12, and 5.

What now? 3

We are now at cave 3.
The neighboring caves are caves 2, 6, 5, and 17.
Brrrr. I feel a breeze, we must be next to a pit.
Whew, what a stench! That is the smell of the Wumpus. It means that we are within two caves of the Wumpus.

What now? Shoot

Figure E.2

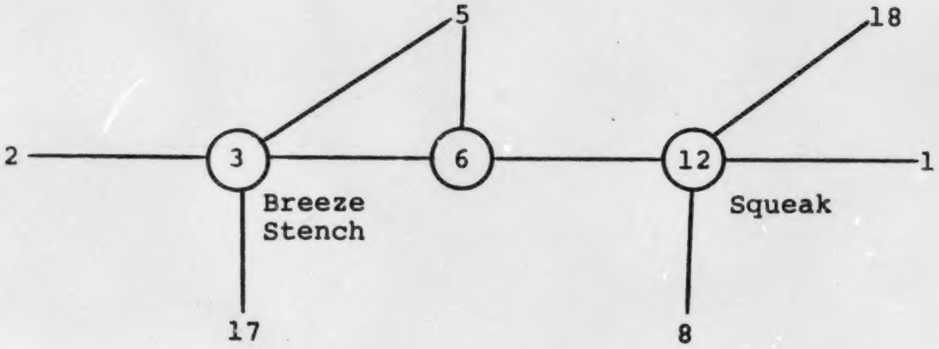

```
Into which cave would you like to shoot?  17

     Cave 17 can not contain the Wumpus because, if the Wumpus
were there we would have smelled him in cave 6 as cave 6 is
within two caves of cave 17.  If we shoot an arrow into cave 17,
it could ricochet back and kill us.

Would you like to shoot into cave 17 anyway? Yes

Poing
Poing
Poing

We are now at cave 3.
The neighboring caves are caves 2, 6, 5, and 17.
Brrrr.  I feel a breeze.
What a stench! The Wumpus is near.

What now? 17

Hummm.
     Jim, we don't need to risk a pit in cave 17 as cave 5 is
safe.  Cave 5 can not contain a pit because, if a pit were there,
we would have felt a breeze in cave 6.  We should try not to
stumble into a pit as it would be fatal.

Would you like to go to cave 17 anyway? No

We are now at cave 3.
The neighboring caves are caves 2, 6, 5, and 17.
Brrrr.
What a stench!

What now? 5

We are now at cave 5.
The neighboring caves are caves 15, 6, 9, and 3.
```

```
Brrrr.

What now? 9
```

Figure E.3

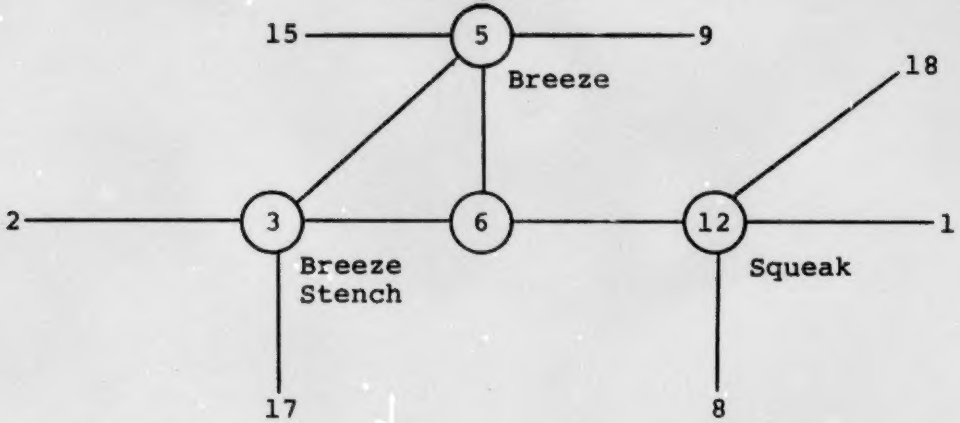

```
Hummm.
    Well, Jim, I think that I would rather risk bats in cave 18
than pits in cave 9.

Would you like to go cave 9 anyway? No

We are now at cave 5.
The neighboring caves are caves 15, 6, 9, and 3.
Brrrr.

What now? Route

What cave would you like to go to? 18
To get to cave 18, Jim, we could move to caves 6, 12, and 18.

Would you like a route to another cave? No

We are now at cave 5.
The neighboring caves are caves 15, 6, 9, and 3.
Brrrr.

What now? 6

We are now at cave 6.
The neighboring caves are caves 3, 12, and 5.

What now? 12

We are now at cave 12.
The neighboring caves are caves 1, 8, 18, and 6.
Squeak.
```

```
What now? 18
```

This game would continue along the same lines. We will shift to a more advanced player.

```
Hello, my name is Wusor VII.  What is your name? Brian CArr

Do I have your name right, Brian Carr? Yes

I believe that we have hunted Wumpii together before. Is that
right? Y

Just a second while I erase my blackboard.

     In this game, Brian, there will be 20 caves all total, three
pits, and three caves with bats.

O.K., now I will draw up a new warren.
We are now at cave 8.
The neighboring caves are caves 6, 9, 13, and 10.
Brrrr.

What now? 6


We are now at cave 6.
The neighboring caves are caves 8, 5, and 0.
Brrrr.

What now? 0
```

Figure E.4

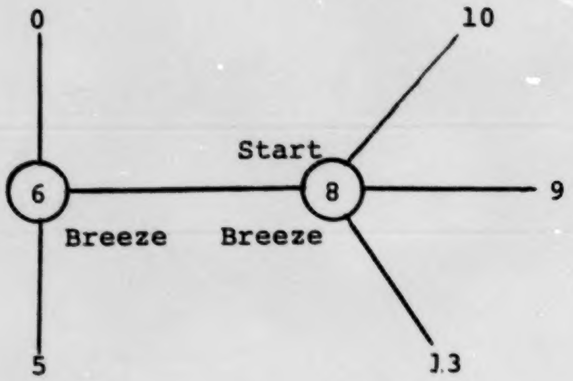

```
Hummm.
    Brian, it isn't necessary to take such large risks with
pits.  It is true that one of caves 13, 9, and 10 contains a
pit because we felt a breeze in cave 8.  However, one of caves
```

```
0 and 5 contains a pit because we felt a breeze in cave 6.
This makes it less likely that cave 9 contains a pit, and so,
Brian, we might want to explore cave 9 instead.

Would you like to go cave 0 anyway? No

We are now at cave 6.
The neighboring caves are caves 8, 5, and 0.
Brrrr.

What now? 8

We are now at cave 8.
The neighboring caves are caves 6, 9, 13, and 10.
Brrrr.

What now? 9

We are now at cave 9.
The neighboring caves are caves 4, 13, 8, and 7.
Brrrr.

What now? 13
```

Figure E.5

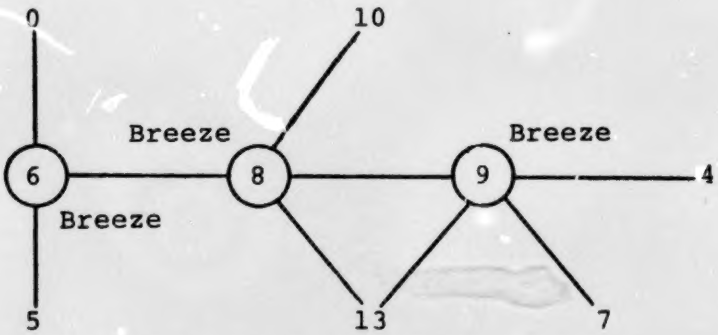

```
Hummm.
    Brian, it isn't necessary to take such large risks with
pits.  One of caves 10 and 13 contain a pit.  Also, one of
caves 7, 4, and 13 contains a pit.  This is multiple evidence
of a pit in cave 13 which makes it quite likely that cave 13
contains a pit.  It is less likely that cave 4 contains a pit.
Hence, Brian, we might want to explore cave 4 instead.

Would you like to go cave 13 anyway? No

We are now at cave 9.
The neighboring caves are caves 4, 13, 8, and 7.
Brrrr.
```

```
What now? Info

What cave would you like to know about? 8

The neighbors of cave 8 are caves 6, 9, 13, and 10.
It has warning "Brrrr.".

Would you like to know about another cave? No

We are now at cave 9.
The neighboring caves are caves 4, 13, 8, and 7.
Brrrr.

What now? 4

We are now at cave 4.
The neighboring caves are caves 7, 3, and 9.
What a stench!

What now? 7
```

Figure E.6

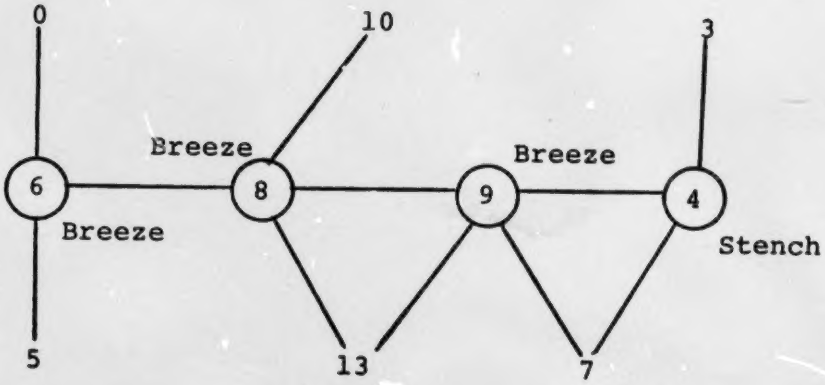

```
Hummm.
     Cave 3 is a very good cave to explore, Brian.  Cave 3 must
be next to the Wumpus because we smelled the Wumpus in cave 4.
We know that cave 4 is not next to the Wumpus, and so cave 4
is two caves away from the Wumpus.  One of caves 2, 16, and 0
must be next to the Wumpus, but we also know that caves 7 and
9 are not next to the Wumpus.  Therefore, cave 3 must be next
to the Wumpus, and, if we visit cave 0, we will gain
information about the location of the Wumpus.

Would you like to go cave 7 anyway? Yes

We are now at cave 7.
The neighboring caves are caves 17, 4, and 9.
Squeak.
```

```
What now? 4

We are now at cave 4.
The neighboring caves are caves 7, 3, and 9.
What a stench!

What now? 3

We are now at cave 3.
The neighboring caves are caves 4, 12, 10, and 18.
What a stench!

What now? 18
```

Figure E.7

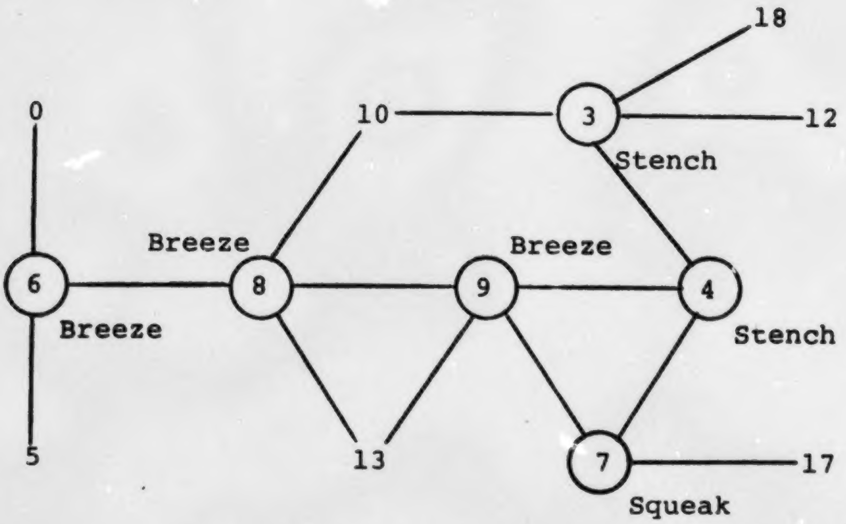

```
     As it is quite possible that cave 18 contains the Wumpus,
Brian, you might want to shoot an arrow into cave 18 before we
visit it.

Would you like to shoot an arrow into cave 18 instead? Yes

Congratulations, Brian, you have shot the Wumpus.
Your record is eight wins and four losses.
Would you like to play again? No

     Well, Brian, have a nice day.  Please look me up next time
you want to go on a Wumpii hunt.

I am returning you to DDT....
```

Results of Experiments With the Synthetic Student

Appendix F

The following are the "results" of the Synthetic Student experiments. The complete protocols are, in fact, much more revealing, but altogether amount to well over a thousand pages of computer output. An editted portion of the last experiment is shown in the following Appendix. This protocol is representative of all of the protocols (except that Syndi normally had a fixed state and did not "learn"). In analyzing the tables, the actual numbers shown are of little importance. The key variable is whether or not the given rule is known (as indicated by an asterisk following the *Familiarity Values*). The first group of experiments were of the Type 1 described earlier.

Experiment with Syndi Possessing Full Knowledge

	Wusor's Student Knowledge Model			Syndi's Student Model		
Rule	Bats	Pits	Wumpus	Bats	Pits	Wumpus
0	6.7830712*	2.78307113*	0.0-	0.0*	0.0*	0.0*
1	386.78307*	278.78307*	305.78307*	0.0*	0.0*	0.0*
2	0.0	0.0	4.7830712*	0.0*	0.0*	0.0*
3	79.78307*	80.78307*	100.78307*	0.0*	0.0*	0.0*
4	27.7830712*	27.7830712*	62.7830715*	0.0*	0.0*	0.0*
5	79.78307*	80.78307*	100.78307*	0.0*	0.0*	0.0*
6	0.0-	0.0-	9.78307116*	0.0*	0.0*	0.0*
7	22.7830712*	27.7830712*	62.7830715*	0.0*	0.0*	0.0*
8	0.0	0.0	31.7830712*	0.0*	0.0*	0.0*
9	0.0	0.0-	0.0-	0.0*	0.0*	0.0*
10	8.78307127*	8.78307127*	4.7830712*	0.0*	0.0*	0.0*
11	18.7830712*	18.7830712*	24.7830712*	0.0*	0.0*	0.0*
12	3.78307113*	2.78307113*	0.0-	0.0*	0.0*	0.0*
13	5.7830712*	3.78307113*	7.7830712*	0.0*	0.0*	0.0*
14	2.78307113*	2.78307113*	0.0	0.0*	0.0*	0.0*
15	0.0	0.0	25.7830712*	0.0*	0.0*	0.0*
19	0.0	0.0	29.7830712*	0.0*	0.0*	0.0*

Experiment with Syndi Possessing Knowledge up to Phase 2

	Wusor's Student Knowledge Model			Syndi's Student Model		
Rule	Bats	Pits	Wumpus	Bats	Pits	Wumpus
0	0.0-	0.0-	0.0-	1.0*	1.0*	1.0*
1	437.78307*	235.78307*	235.78307*	11.0*	11.0*	11.0*
2	0.0	0.0	0.0-	0.0	0.0	1.0*
3	27.7830712*	33.7830715*	8.78307127*	1.0*	1.0*	1.0*
4	1.0-	1.0-	3.78307113*	0.0	0.0	0.0
5	27.7830712*	33.7830715*	8.78307127*	1.0*	1.0*	1.0*
6	0.0	0.0	0.0-	0.0	0.0	0.0
7	1.0-	1.0-	3.78307113*	0.0	0.0	0.0
8	0.0	0.0	0.0-	0.0	0.0	0.0
9	0.0	0.0	0.0-	0.0	0.0	0.0
10	0.0	0.0	0.0-	0.0	0.0	0.0
11	1.0-	1.0-	0.0-	0.0	0.0	0.0
12	0.0	0.0	0.0	0.0	0.0	0.0
13	0.0	0.0	0.0	0.0	0.0	0.0
14	0.0	0.0	0.0	0.0	0.0	0.0
15	0.0	0.0	0.0-	0.0	0.0	0.0
19	0.0	0.0	2.78307113*	0.0	0.0	1.0*

Experiment with Syndi Possessing Knowledge up to Phase 3

	Wusor's Student Knowledge Model			Syndi's Student Model		
Rule	Bats	Pits	Wumpus	Bats	Pits	Wumpus
0	2.78307113*	2.78307113*	0.0-	1.0*	1.0*	0.0
1	130.78307*	92.78307*	95.78307*	34.0*	34.0*	34.0*
2	0.0	0.0	4.7830712*	0.0	0.0	2.0*
3	21.7830712*	24.7830712*	21.7830712*	1.0*	1.0*	5.0*
4	8.78307127*	8.78307127*	15.7830713*	2.0*	2.0*	4.0*
5	21.7830712*	24.7830712*	21.7830712*	1.0*	1.0*	5.0*
6	0.0-	0.0-	4.7830712*	0.0	0.0	0.0
7	3.78307113*	8.78307127*	15.7830713*	2.0*	2.0*	4.0*
8	0.0	0.0	8.78307127*	0.0	0.0	0.0
9	0.0	0.0-	0.0-	0.0	0.0	0.0
10	3.78307113*	3.78307113*	1.0-	2.0*	2.0*	0.0
11	6.7830712*	6.7830712*	8.78307127*	1.0*	1.0*	0.0
12	1.0-	1.0-	0.0-	0.0	0.0	0.0
13	1.0-	1.0-	0.0-	0.0	0.0	0.0
14	0.0	0.0	0.0	0.0	0.0	0.0
15	0.0	0.0	8.78307127*	0.0	0.0	0.0
19	0.0	0.0	8.78307127*	0.0	0.0	3.0*

The next set of experiments were of the Type 2 with Wusor initially assuming that Syndi had a complete knowledge of the game. Because of the nature of the Phases, once the player is no longer presumed to know a simple rule (like rules 3 and 5), the higher rules can no longer be degraded. In any case, the player was reduced to the appropriate phase, and advice would only be given at the appropriate level.

Experiment With Wusor Degrading his Student Model

	Wusor's Student Knowledge Model			Syndi's Student Model		
Rule	Bats	Pits	Wumpus	Bats	Pits	Wumpus
0	0.84-	1.0*+	0.0	1.0*	1.0*	1.0*
1	120.0*+	112.0*+	113.0*+	1.0*	1.0*	1.0*
2	0.0	0.0	0.0-	0.0	0.0	1.0*
3	0.34-+	0.34-+	0.34-+	0.0	0.0	0.0
4	3.0*+	2.0*+	3.0*+	0.0	0.0	0.0
5	0.34-+	0.0-+	0.34-+	0.0	0.0	0.0
6	0.0-	0.0	0.0	0.0	0.0	0.0
7	3.0*+	2.0*+	3.0*+	0.0	0.0	0.0
8	0.0	0.0	2.0*+	0.0	0.0	0.0
9	0.0	0.0	1.0*	0.0	0.0	0.0
10	1.0*+	1.0*+	0.0	0.0	0.0	0.0
11	3.0*+	2.0*+	0.0 +	0.0	0.0	0.0
12	1.0*+	1.0*+	0.0 +	0.0	0.0	0.0
13	1.0*+	0.0 +	1.0*+	0.0	0.0	0.0
14	1.0*+	0.34 +	0.0	0.0	0.0	0.0
15	0.0	0.0	0.0 +	0.0	0.0	0.0
19	0.0	0.0	0.0-+	0.0	0.0	0.0

The last set of experiments were of the Type 3, where Syndi started with a minimal knowledge of the game (rules 0, 1, and 2) and "learned" rules as they were explained to her.

Experiment with Syndi Learning

	Wusor's Student Knowledge Model			Syndi's Student Model		
Rule	Bats	Pits	Wumpus	Bats	Pits	Wumpus
0	10.0*+	4.0*	0.0-	3.0*	1.0*	1.0*
1	517.0*+	496.0*+	423.0*+	6.0*	4.0*	29.0*
2	0.0	0.0	7.0*+	0.0	0.0	9.0*
3	118.0*	118.0*+	68.0*+	1.0*	1.0*	29.0*
4	50.0*+	51.0*	52.0*+	3.0*	3.0*	27.0*
5	118.0*	118.0*+	68.0*+	1.0*	1.0*	29.0*
6	0.0-	0.0-	2.0*	0.0	0.0	0.0
7	33.0*+	51.0*	52.0*+	4.0*	3.0*	32.0*
8	0.0	0.0	35.0*+	0.0	0.0	32.0*
9	0.0	0.0-	0.0-	0.0	0.0	0.0
10	13.0*+	13.0*	6.0*+	1.0*	1.0*	11.0*
11	38.0*+	39.0*	28.0*+	1.0*	1.0*	4.0*
12	0.0-	0.0-	0.0-	0.0	0.0	0.0
13	10.0*	3.0*+	4.0*	1.0*	1.0*	0.0
14	0.0	0.0	0.0	0.0	0.0	0.0
15	0.0	0.0	31.0*	0.0	0.0	0.0
19	0.0	0.0	22.0*+	0.0	0.0	0.0

Scenario Between Syndi Learn and Wusor

Appendix G

What follows is a typical protocol with the Synthetic Student and Wusor. At various points each had inserted their current student model values into the file, but these were deleted to save space. Also, between every move both Wusor and Syndi would output their analyses of each move, but most of these have also been deleted to conserve space (the remaining are shown in italics). The complete protocols from all the experiments are available through the author.

Comments made by Wusor which would not normally be displayed for the student are enclosed in "***". They constitute a trace of Wusor's reasoning and make Wusor more believable to a human teacher as well as serving as an invaluable debugging aid. Syndi's comments are much less frequent and are enclosed in "###". Syndi's replies are all listed in upper case characters and underlined.

```
Hello, my name is Wusor VIII.  What is your name?  (Please
type in your first and last name.  Follow all responses with
a space.) SYNDI LEARN

Do I have your name right, Syndi Learn? YES

I believe that we have hunted Wumpii together before.  Is
that right? YES

      In this game, Syndi, there will be 20 caves all
total, one pit, and one cave with bats.

Just a second while I erase my blackboard.
*** Time is 190. ***
O.K., now I will draw up a new warren.
*** I am altering the game to create a situation appropriate
for a player of phase 0. ***
We are now at cave 11.
The neighboring caves are caves 14, 1, 9, and 19.

What now?
```

```
### The following is my estimation of the moves.
Cave     Cost          Gain          Index
  19     0.198270723   0.80131574    4.0415233
  9      0.198270723   0.80131574    4.0415233
  1      0.198270723   0.80131574    4.0415233
  14     0.198270723   0.80131574    4.0415233
       A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave     Cost         Gain         Index
  19     0.0          1.0          10000.0
  9      0.0          1.0          10000.0
  1      0.0          1.0          10000.0
  14     0.0          1.0          10000.0
14

We are now at cave 14.
The neighboring caves are caves 16, 5, 19, and 11.
Brrrr.  I feel a breeze!  We must be next to a pit.

What now?

### The following is my estimation of the moves.
Cave     Cost          Gain          Index
  9      0.198270723   0.79707436    4.0201313
  1      0.198270723   0.79707436    4.0201313
  19     0.198270723   0.79707436    4.0201313
  5      0.198270723   0.79707436    4.0201313
  16     0.198270723   0.79707436    4.0201313
       A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave     Cost         Gain         Index
  5      0.5          0.5          1.0
  16     0.5          0.5          1.0
  9      0.0          1.0          10000.0
  1      0.0          1.0          10000.0
  19     0.0          1.0          10000.0
16

      Syndi, why not go back to cave 11 where we didn't
get any warnings at all.  From there we can safely go to
cave 1 without risking the pit.  We should try not to
stumble into the pit as it would be fatal.
*** I have used: (rule 1 for the Wumpus), (rule 1 for the
pit), and (rule 1 for the bats) in explaining combination
rule 8. ***
Would you like to go to cave 16 anyway? NO

We are now at cave 14.
The neighboring caves are caves 16, 5, 19, and 11.
Brrrr.  I feel a breeze.

What now? 11
```

```
*** I have stopped initializing the player's knowledge model
for the bats because he seems to be learning rule 1. ***
*** I am advancing the student to phase 1 for the bats as he
has mastered rule 1. ***
*** I have stopped initializing the player's knowledge model
for the pit because he seems to be learning rule 1. ***
*** I am advancing the student to phase 1 for the pit as he
has mastered rule 1. ***
*** I have stopped initializing the player's knowledge model
for the Wumpus because he seems to be learning rule 1. ***
*** I am advancing the student to phase 1 for the Wumpus as
he has mastered rule 1. ***
We are now at cave 11.
The neighboring caves are caves 14, 1, 9, and 19.

What now? 1

*** I am upgrading the player's receptivity because his move
to cave 1 is a move that I told him was good very recently.
***
We are now at cave 1.
The neighboring caves are caves 0, 15, and 11.
Whew, what a stench!  That is the smell of the Wumpus.  It
means that we are within two caves of the Wumpus.

What now?

### The following is my estimation of the moves.
Cave    Cost         Gain         Index
  5     0.198270723  0.79283295   3.9987394
  16    0.198270723  0.79283295   3.9987394
  19    0.198270723  0.79283295   3.9987394
  9     0.198270723  0.79283295   3.9987394
  15*   0.1          2.792833     27.9283297
  0*    0.1          2.792833     27.9283297
      A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost         Gain         Index
  5     0.5          0.5          1.0
  16    0.5          0.5          1.0
  19    0.0          1.0          10000.0
  9     0.0          1.0          10000.0
  15    0.0          1.1          11000.0
  0     0.0          1.1          11000.0
SHOOT

Into which cave would you like to shoot? 0

      Cave 0 can not contain the Wumpus because, if the
Wumpus were there, we would have smelled the Wumpus in cave
11 as cave 0 is within two caves of cave 11, and, if we
shoot an arrow into cave 0, it could ricochet back and kill
us.
```

```
*** I have used: (rules 3 and 5 for the Wumpus) in
explaining combination rule 9. ***
### I have just "learned" rule 3 for the Wumpus. ###
### I have just "learned" rule 5 for the Wumpus. ###
Would you like to shoot into cave 0 anyway? NO

We are now at cave 1.
The neighboring caves are caves 0, 15, and 11.
What a stench!  The Wumpus is near.

What now? SHOOT

Into which cave would you like to shoot? 15

Poing
You have just shot yourself.
Would you like to take back your last move? YES

We are now at cave 1.
The neighboring caves are caves 0, 15, and 11.
What a stench!

What now? 15

We are now at cave 15.
The neighboring caves are caves 1, 6, 3, and 10.
Squeak.  I hear bats, they must be in one of the neighboring
caves.
What a stench!

What now?

### The following is my estimation of the moves.
Cave    Cost           Gain          Index
  5     0.198270723    0.78859156    3.9773475
  16    0.198270723    0.78859156    3.9773475
  19    0.198270723    0.78859156    3.9773475
  9     0.198270723    0.78859156    3.9773475
  0     0.198270723    0.78859156    3.9773475
  10"   0.1            2.78859156    27.8859157
  3"    0.1            2.78859156    27.8859157
  6"    0.1            2.78859156    27.8859157
     A """ after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost           Gain          Index
  5     0.5            0.5           1.0
  16    0.5            0.5           1.0
  10    0.227228693    0.7828529     3.4452203
  3     0.227228693    0.7828529     3.4452203
  6     0.227228693    0.7828529     3.4452203
  19    0.0            1.0           10000.0
  9     0.0            1.0           10000.0
  0     0.0            1.1           11000.0
```

```
SHOOT

Into which cave would you like to shoot? 6

Congratulations, Syndi, you have shot the Wumpus.
Your record is one wins and one losses.
Would you like to play again? YES

      In this game, Syndi, there will be 20 caves all
total, two pits, and two caves with bats.

Just a second while I erase my blackboard.
*** Time is 149. ***
O.K., now I will draw up a new warren.
*** I am altering the game to create a situation appropriate
for a player of phase 0. ***

We are row at cave 3.
The neighboring caves are caves 7, 0, and 15.

What now? 7

We are now at cave 7.
The neighboring caves are caves 6, 3, 13, and 18.
Squeak.  I hear bats.
Brrrr.

What now?

### The following is my estimation of the moves.
Cave      Cost          Gain          Index
  0       0.114270098   0.88430289    7.7387077
  15      0.114270098   0.88430289    7.7387077
  18      0.114270098   0.88430289    7.7387077
  13      0.114270098   0.88430289    7.7387077
  6       0.114270098   0.88430289    7.7387077
      A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave      Cost          Gain          Index
  18      0.36856813    0.6279084     1.70364266
  13      0.36856813    0.6279084     1.70364266
  6       0.36856813    0.6279084     1.70364266
  0       0.0           1.0           10000.0
  15      0.0           1.0           10000.0
6

Hummmm.
*** Better moves which I have not mentioned to the student
recently are caves 15 and 0. ***
      Syndi, we don't need to risk bats and pits in cave 6
as cave 15 is safe.  Cave 15 can not contain a pit because,
if a pit were there, we would have felt a breeze in cave 3.
We should avoid pits as they are fatal.  Likewise, cave 15
```

```
is safe from the other dangers.
*** I have used: (rules 3 and 5 for pits) in explaining
combination rule 8. ***
### I have just "learned" rule 3 for pits. ###
### I am transferring my knowledge of rule 3 for pits to
bats. ###
### I have just "learned" rule 5 for pits. ###
### I am transferring my knowledge of rule 5 for pits to
bats. ###
Would you like to go to cave 6 anyway? NO

We are now at cave 7.
The neighboring caves are caves 6, 3, 13, and 18.
Squeak.
Brrrr.

What now? 3

We are now at cave 3.
The neighboring caves are caves 7, 15, and 8.

What now? 15

*** I am upgrading the player's receptivity because his move
to cave 15 is a move that I told him was good very recently.
***
We are now at cave 15.
The neighboring caves are caves 1, 4, 3, and 8.

What now?

### The following is my estimation of the moves.
Cave     Cost          Gain          Index
  18     0.114270098   0.87938217    7.6956455
  13     0.114270098   0.87938217    7.6956455
  6      0.114270098   0.87938217    7.6956455
  0      0.0           1.0           10000.0
  8      0.0           1.0           10000.0
  4      0.0           1.0           10000.0
  1      0.0           1.0           10000.0
     A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave     Cost          Gain          Index
  18     0.36856813    0.61575847    1.67067747
  13     0.36856813    0.61575847    1.67067747
  6      0.36856813    0.61575847    1.67067747
  0      0.0           1.0           10000.0
  8      0.0           1.0           10000.0
  4      0.0           1.0           10000.0
  1      0.0           1.0           10000.0
1

We are now at cave 1.
```

```
The neighboring caves are caves 15, 2, 14, and 9.
Squeak.
What a stench!

What now? 15

We are now at cave 15.
The neighboring caves are caves 1, 4, 3, and 8.

What now? 4

*** Moves which the student seems to have correctly
identified as worse are caves 2, 14, and 9. ***
*** According to combination rule 0, I am marking rules 3
and 5 for bats because his move is better than cave 2. ***
We are now at cave 4.
The neighboring caves are caves 15, 11, 18, and 2.
Squeak.
Brrrr.

What now? 15

We are now at cave 15.
The neighboring caves are caves 1, 4, 3, and 8.

What now? 8

*** Moves which the student seems to have correctly
identified as worse are caves 2, 11, and 18. ***
*** According to combination rule 0, I am marking rules 3
and 5 for bats because his move is better than cave 2. ***
*** I am advancing the student to phase 2 for bats as he has
mastered rules 3 and 5. ***
*** According to combination rule 0, I am marking rules 3
and 5 for pits because his move is better than cave 11.
***
*** I am advancing the student to phase 2 for pits as he has
mastered rules 3 and 5. ***
We are now at cave 8.
The neighboring caves are caves 15, 19, 13, and 2.
Brrrr.

What now? 2

We are now at cave 2.
The neighboring caves are caves 9, 1, 8, and 4.
What a stench!

What Now? 9

We are now at cave 9.
The neighboring caves are caves 10, 2, 1, and 5.
What a stench!
What now? 1

We are now at cave 1.
The neighboring caves are caves 15, 2, 14, and 9.
Squeak.
What a stench!

What now? 15

We are now at cave 15.
The neighboring caves are caves 1, 4, 3, and 8.

What now? 3

We are now at cave 3.
The neighboring caves are caves 7, 15, and 8.

What now? 8

*** Moves which the student seems to have correctly
identified as worse are caves 10, 5, 14, 19, 11, 13,
and 18. ***
*** According to combination rule 8, I am marking rules 3
and 5 for the Wumpus because his move is better than cave 10.
***
*** According to combination rule 8, I am marking rules 3
and 5 for bats because his move is better than cave 14.
***
*** According to combination rule 8, I am marking rules 3
and 5 for pits becausa his move is better than cave 19.
***
We are now at cave 8.
The neighboring caves are caves 11, 16, and 3.
What a stench!

What now?

### The following is my estimation of the moves.
Cave     Cost          Gain         Index
  18     0.114270098  0.84985783   7.4372723
  6      0.114270098  0.84985783   7.4372723
  19     0.100000001  0.9          8.99999988
  13     0.100000001  0.9          8.99999988
  5*     0.1          2.9          29.0000002
  10*    0.1          2.9          29.0000002
  14     0.0158556625 0.9442866    59.555159
  16     0.0          1.0          10000.0
  11     0.0          1.0          10000.0
      A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave     Cost         Gain          Index
  18     1.0          0.0          0.0
  19     0.6          0.6          1.0
  13    0.5           0.5          1.0
  14    0.158556588   0.442864964  2.79310346
  5"    0.16665       3.416685     20.5021603
  10"   0.16365       3.416685     20.5021603
  6     0.0           1.0          10000.0
  11    0.0           1.0          10000.0
  16    0.0           1.1          11000.0
11

Hummmm.
*** Better moves which I have not mentioned to the student
recently are cave 16. ***
We are now at cave 11.
The neighboring caves are caves 8, 4, 17, and 18.
Squeak.
Brrrr.

What now? 8

We are now at cave 8.
The neighboring caves are caves 11, 16, and 3.
What a stench!

What now? 16

*** Moves which the student seems to have correctly
identified as worse are caves 17 and 18. ***
We are now at cave 16.
The neighboring caves are caves 13, 5, 18, and 8.
Brrrr.
What a stench!

What now?

### The following is my estimation of the moves.
Cave    Cost          Gain          Index
  8     0.114270098   0.8400164     7.3511479
  18    0.114270098   0.8400164     7.3511479
  17    0.114270098   0.8400164     7.3511479
  19    0.100000001   0.9           8.99999988
  13    0.100000001   0.9           8.99999988
  10"   0.1           2.9           29.0000002
  5"    0.1           2.9           29.0000002
  14    0.0168556626  0.93336156    58.865503
      A """ after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost          Gain          Index
  18    1.0           0.0           0.0
  13    1.0           0.0           0.0
  14    0.158556588   0.333515584   2.1034483
  10"   0.16665       3.416685      20.5021603
  5"    0.16665       3.416685      20.5021603
  19    0.0           1.0           10000.0
```

```
   6     0.0           1.0           10000.0
   17    0.0           1.0           10000.0
8

We are now at cave 8.
The neighboring caves are caves 11, 16, and 3.
What a stench!

What now? 3

We are now at cave 3.
The neighboring caves are caves 7, 15, and 8.

What now? 15

We are now at cave 15.
The neighboring caves are caves 1, 4, 3, and 8.

What now? 1

We are now at cave 1.
The neighboring caves are caves 15, 2, 14, and 9.
Squeak.
What a stench!

What now? 14

*** Moves which the student seems to have correctly
identified as worse are caves 13 and 18. ***
Hummmm.
*** Better moves which I have not mentioned to the student
recently are caves 19, 17, 10, and 5. ***
Bon Voyage!
We are now at cave 8.
The neighboring caves are caves 15, 19, 13, and 2.
Brrrr.

What now?

### The following is my estimation of the moves.
Cave    Cost            Gain           Index
  14    0.158566588   0.333616584   2.1034483
  6     0.114270098   0.8400164     7.3611479
  17    0.114270098   0.8400164     7.3611479
  18    0.114270098   0.8400164     7.3611479
  13    0.100000001   0.9           8.99999988
  19    0.100000001   0.9           8.99999988
  5*    0.1           2.9           29.0000002
  10*   0.1           2.9           29.0000002
      A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost          Gain          Index
  18    1.0           0.0           0.0
```

```
    13     1.0           0.0           0.0
    14     0.158566688   0.333516584   2.1034483
    5"     0.16665       3.416685      20.5021603
    10"    0.16665       3.416685      20.5021603
    6      0.0           1.0           10000.0
    17     0.0           1.0           10000.0
    19     0.0           1.0           10000.0
  2

  We are now at cave 2.
  The neighboring caves are caves 9, 1, 8, and 4.
  What a stench!

  What now? 9

  We are now at cave 9.
  The neighboring caves are caves 10, 2, 1, and 5.
  What a stench!

  What now? SHOOT

  Into which cave would you like to shoot? 10

  Poing
  We are now at cave 9.
  The neighboring caves are caves 10, 2, 1, and 5.
  What a stench!

  What now?

  ### The following is my estimation of the moves.
  Cave     Cost          Gain          Index
    14     0.158566688   0.333516584   2.1034483
    6      0.114270098   0.8400164     7.3511479
    17     0.114270098   0.8400164     7.3511479
    18     0.114270098   0.8400164     7.3511479
    13     0.100000001   0.9           8.99999988
    19     0.100000001   0.9           8.99999988
    5"     0.1           2.9           29.0000002
    10     0.0           1.0           10000.0
        A """ after the cave number indicates that these
  figures reflect shooting into the cave before visiting it.
  Cave     Cost          Gain          Index
    18     1.0           0.0           0.0
    13     1.0           0.0           0.0
    14     0.158566688   0.333516584   2.1034483
    6      0.0           1.0           10000.0
    17     0.0           1.0           10000.0
    19     0.0           1.0           10000.0
    10     0.0           1.0           10000.0
    5"     0.0           5.1           51000.0
  10
```

```
Hummmm.
*** Better moves which I have not mentioned to the student
recently are cave 5. ***
We are now at cave 10.
The neighboring caves are caves 17, 9, 19, and 16.
What a stench!

What now? 17

Hummmm.
*** Better moves which I have not mentioned to the student
recently are cave 5. ***
We are now at cave 17.
The neighboring caves are caves 6, 10, 12, and 11.
What a stench!

What now? 6

*** I am degrading the player's receptivity because his move
 to cave 6 is a move that I told him was bad very recently.
***
We are now at cave 6.
The neighboring caves are caves 7, 12, 17, and 13.
Brrrr.
What a stench!

What now? 12

Hummmm.
*** Better moves which I have not mentioned to the student
recently are cave 5. ***
We are now at cave 12.
The neighboring caves are caves 5, 6, 14, and 17.
Squeak.
What a stench!

What now? 17

We are now at cave 17.
The neighboring caves are caves 6, 10, 12, and 11.
What a stench!

What now? 10

We are now at cave 10.
The neighboring caves are caves 17, 9, 19, and 16.
What a stench!

What now? 19

*** Moves which the student seems to have correctly
identified as worse are caves 14 and 13. ***
*** According to combination rule 0, I am marking rules 3
```

```
and 5 for bats because his move is better than cave 14.
***
*** According to combination rule 0, I am marking rules 3
and 5 for pits because his move is better than cave 13.
***
Hummmm.
*** Better moves which I have not mentioned to the student
recently are cave 5. ***
We are now at cave 19.
The neighboring caves are caves 14, 8, 18, and 10.
Squeak.
Brrrr.
What a stench!

What now?

### The following is my estimation of the moves.
Cave    Cost         Gain         Index
  14    0.158555588  0.060142154  0.37931034
  18    0.114270098  0.81541279   7.135837
  13    0.100000001  0.9          8.99999988
  5"    0.1          2.9          29.0000002
     A """ after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost         Gain         Index
  13    1.0          0.0          0.0
  18    1.0          0.0          0.0
  14    0.158555588  0.060142154  0.37931034
  5"    0.0          5.1          51000.0
10

We are now at cave 10.
The neighboring caves are caves 17, 9, 19, and 16.
What a stench!

What now? 16

We are now at cave 16.
The neighboring caves are caves 13, 5, 10, and 0.
Brrrr.
What a stench!

What now? SHOOT

Into which cave would you like to shoot? 5

Congratulations, Syndi, you have shot the Wumpus.
Your record is two wins and one losses.
Would you like to play again? YES

     In this game, Syndi, there will be 20 caves all
total, two pits, and two caves with bats.
```

```
Just a second while I erase my blackboard.
*** Time is 208. ***
O.K., now I will draw up a new warren.
*** I am altering the game to create a situation appropriate
for a player of phase 8. ***

We are now at cave 5.
The neighboring caves are caves 8, 12, 3, and 19.

What now? 8

We are now at cave 8.
The neighboring caves are caves 5, 4, 6, and 7.
Squeak.
Brrrr.
What a stench!

What now?

### The following is my estimation of the moves.
Cave    Cost          Gain          Index
  7     0.114270098   0.88430289    7.7387077
  6     0.114270098   0.88430289    7.7387077
  4     0.114270098   0.88430289    7.7387077
  19    0.0           1.0           10000.0
  12    0.0           1.0           10000.0
  3     0.0           1.0           10000.0
     A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost          Gain          Index
  7     0.36856813    0.67236283    1.82422945
  6     0.36856813    0.67236283    1.82422945
  4     0.36856813    0.67236283    1.82422945
  19    0.0           1.0           10000.0
  12    0.0           1.0           10000.0
  3     0.0           1.0           10000.0
5

We are now at cave 5.
The neighboring caves are caves 8, 3, 12, and 19.

What now? 3

*** Moves which the student seems to have correctly
identified as worse are caves 4, 6, and 7. ***
We are now at cave 3.
The neighboring caves are caves 18, 9, and 5.

What now? 18

We are now at cave 18.
The neighboring caves are caves 3, 17, and 14.
```

```
What now? 17

We are now at cave 17.
The neighboring caves are caves 15, 18, 11, and 14.
Brrrr.

What now?

### The following is my estimation of the moves.
Cave     Cost         Gain         Index
  7      0.114270098  0.86954073   7.6095211
  6      0.114270098  0.86954073   7.6095211
  4      0.114270098  0.86954073   7.6095211
  11     0.100000001  0.9          8.99999988
  15     0.100000001  0.9          8.99999988
  19     0.0          1.0          10000.0
  12     0.0          1.0          10000.0
  9      0.0          1.0          10000.0
  14     0.0          1.0          10000.0
      A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave     Cost         Gain         Index
  11     0.5          0.5          1.0
  15     0.5          0.5          1.0
  7      0.36856813   0.63590305   1.7253338
  6      0.36856813   0.63590305   1.7253338
  4      0.36856813   0.63590305   1.7253338
  19     0.0          1.0          10000.0
  12     0.0          1.0          10000.0
  9      0.0          1.0          10000.0
  14     0.0          1.0          10000.0
14

We are now at cave 14.
The neighboring caves are caves 8, 11, 18, and 17.
Squeak.
Brrrr.

What now? 18

We are now at cave 18.
The neighboring caves are caves 3, 17, and 14.

What now? 3

We are now at cave 3.
The neighboring caves are caves 18, 9, and 5.

What now? 9

*** Moves which the student seems to have correctly
identified as worse are caves 8 and 11. ***
*** According to combination rule 8, I am marking rules 3
```

```
and 5 for bats because his move is better than cave 8. ***
*** According to combination rule 0, I am marking rules 3
and 5 for pits because his move is better than cave 11.
***
We are now at cave 9.
The neighboring caves are caves 3, 6, 10, and 2.
What a stench!

What now? 6

We are now at cave 6.
The neighboring caves are caves 8, 9, 11, and 1.
Brrrr.
What a stench!

What now? 9

We are now at cave 9.
The neighboring caves are caves 3, 6, 10, and 2.
What a stench!

What now? 10

*** Moves which the student seems to have correctly
identified as worse are caves 1, 4, 7, and 11. ***
*** According to combination rule 0, I am marking rules 3,
5, and 19 for the Wumpus because his move is better than
cave 1. ***
*** I am advancing the student to phase 2 for the Wumpus as
he has mastered rules 3, 5, and 19. ***
*** According to combination rule 0, I am marking rules 3
and 5 for pits because his move is better than cave 11.
***
We are now at cave 10.
The neighboring caves are caves 9, 2, 15, and 13.
What a stench!

What now? 2

We are now at cave 2.
The neighboring caves are caves 10, 0, 16, and 9.
Squeak.
What a stench!

What now? 10

We are now at cave 10.
The neighboring caves are caves 9, 2, 15, and 13.
What a stench!

What now? 15

*** Moves which the student seems to have correctly
```

```
identified as worse are caves 16, 13, and 8. ***
*** According to combination rule 8, I am marking rules 3
and 5 for the Wumpus because his move is better than cave 16.
***
*** According to combination rule 8, I am marking rules 3
and 5 for bats because his move is better than cave 8. ***
We are now at cave 15.
The neighboring caves are caves 17, 19, and 10.

What now?

### The following is my estimation of the moves.
Cave    Cost         Gain        Index
  7     0.114270098  0.8400164   7.3511479
  4     0.114270098  0.8400164   7.3511479
  11    0.100000001  0.9         8.99999988
  1*    0.1          2.81        28.1
  16*   0.1          2.8400164   28.400164
  0     0.0158556625 0.93335155  58.865503
  12    0.0          1.0         10000.0
  13    0.0          1.0         10000.0
  19    0.0          1.0         10000.0
      A "*" after the cave number indicates that these
figures reflect shooting into the cave before visiting it.
Cave    Cost         Gain        Index
  11    1.0          0.0         0.0
  7     0.53963915   0.358378894 0.6641084
  4     0.53963915   0.358378894 0.6641084
  0     0.158556588  0.333515584 2.1034483
  1*    0.161094997  3.5061289   21.764356
  16*   0.127765     4.0427918   31.642405
  12    0.0          1.0         10000.0
  19    0.0          1.0         10000.0
  13    0.0          1.1         11000.0
19

Hummmm.
*** Better moves which I have not mentioned to the student
recently are cave 13. ***
        Cave 13 is a very good cave to explore, Syndi, as we
know that one of caves 2 and 13 must be next to the Wumpus.
Cave 10 must be exactly two caves away from the Wumpus
because we smelled the Wumpus in cave 10.  This means that
cave 10 is within two caves of the Wumpus.  We also know
that cave 10 is not next to the Wumpus.  Hence cave 10 is
exactly two caves away from the Wumpus.  Therefore one of
caves 9, 2, 15, and 13 must be next to the Wumpus, but
neither of caves 9 and 15 are next to the Wumpus.  This
means that one of caves 2 and 13 must be next to the Wumpus,
and so if we visit cave 13 we will gain information about
the location of the Wumpus.
*** I have used: (rules 8, 4, and 7 for the Wumpus) in
explaining combination rule 5. ***
```

```
### I have just "learned" rule 8 for the Wumpus. ###
### I have just "learned" rule 4 for the Wumpus. ###
### I have just "learned" rule 7 for the Wumpus. ###
Would you like to go to cave 19 anyway?
```

This experiment continued on for many more games.

Bibliography

Artificial Intelligence
Computer Aided Instruction

Barr, A., M. Beard and R. C. Atkinson, The Computer as a Tutorial Laboratory: The Stanford BIP Project, Technical Report No. 260, Psychology and Education Series, Institute for Mathematical Studies in the Social Sciences, August, 1975.

Bork, A.M., "Effective Computer Use in Physics Education", American Journal of Physics, January 1975.

Brown, J.S., "Uses of Artificial Intelligence and Advanced Computer Technology in Education", to appear in Proceedings of a Conference on the Impact of Computers on Education, A Ten Year Forecast, Washington, D.C., forthcoming 1976.

Brown, J.S., Bell, A. & Zydbel, F. Steps Toward an Artificially Intelligent CAI System, University of California at Irvine ICS Technical Report No. 35, July 1973.

Brown, J.S. & Burton, R., "Multiple Representations of Knowledge for Tutorial Reasoning", in D. Bobrow & A. Collins (Eds.), Representations and Understanding Studies in Cognitive Science, New York: Academic Press, 1975.

Brown, J.S., Burton, R. & Bell, A., "SOPHIE: A Step Toward Creating a Reactive Learning Environment", International Journal of Man-Machine Studies, Vol. 7, 1975, pp. 675-696.

Brown, J.S., Burton, R.R. & Zdybel, F. "A Model-Driven Question-Answering System for Mixed-Initiative Computer-Assisted Instruction", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-3, No. 3, May 1973, pp. 248-257.

Burton, R. & J. S. Brown, "A Tutoring and Student Modelling Paradigm for Gaming Environments", in R. Coleman and P. Lorton, Jr. (Eds.), Computer Science and Education, (Advance Proceedings of the Association for Computer Machinery Special Interest Groups on Computer Science Education and Computer Uses in Education Joint Symposium, Anaheim, CA), SIGCSE Bulletin, Vol. 8, No. 1, Feb. 1976, pp. 236-246.

Carbonell, J., "AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction", IEEE Transactions on Man-Machine Systems, Vol. MMS-11, No. 4, December 1970.

Carbonell, J., Mixed-Initiative Man-Computer Instructional Dialogues, MIT Ph.D. Thesis, June 1970.

Carbonell, J. & Collins, A., "Natural Semantics in Artificial Intelligence", in Proceedings of the Third International Joint Conference on Artificial Intelligence, Stanford University, 1973, pp. 344-351.

Carr, B. & Goldstein I., Overlays: a Theory of Modelling for Computer Aided Instruction, MIT Artificial Intelligence Laboratory Memo No. 406, February 1977.

Collins, A., Comparison of Two Teaching Strategies in Computer-Assisted Instruction, BBN Report No.2885, 1974.

Collins, A., "Education and Understanding", in D. Klahr (Ed.), Cognition and Instruction, Hillsdale, NJ: Erlbaum Associates, 1975.

Collins, A., Processes in Acquiring Knowledge, BBN Report No. 3231, January, 1976.

Collins, A., Adams, M.J. & Pew, R.W., The Effectiveness of an Interactive Map Display in Tutoring Geography, BBN Report No. 3346, August, 1976.

Collins, A. & Grignetti, M., Intelligent CAI, BBN Report No. 3181, 1975.

Collins, A., Passafiume, J., Gould, L. & Carbonell, J., Improving Interactive Capabilities in Computer-Assisted Instruction, BBN Report No. 2631, 1973.

Collins, A. & Warnock, E., Semantic Networks, BBN Report No. 2833, Job No. 11489, May 1974.

Collins, A., Warnock, E., Aiello, N. & Miller, M., "Reasoning from Incomplete Knowledge", in D. Bobrow and A. Collins (Eds.), Representation and Understanding Studies in Cognitive Science, New York: Academic Press, 1975.

Collins, A., Warnock, E. & Passafiume, J., "Analysis and Synthesis of Tutorial Dialogues", in G. Bower (Ed.), The Psychology of Learning and Motivation, Vol. 9, New York: Academic Press, 1975.

Edelson, G.D., An Artificial Intelligence Analysis of Algebra Problem Solving, MIT B.S. Thesis, May 1976.

Gallenson, L., An Approach to Providing a User Interface for Military Computer-Aided Instruction, Information Sciences Institute RR-75-43, November 1975.

Goldstein, I. P. The Computer as Coach: an Athletic Paradigm for Intellectual Education, MIT Artificial Intelligence Laboratory, Memo 389, February 1977.

Grignetti, M., Hausmann, C. & Gould, L., "An Intelligent' On-Line Assistant and Tutor -- NLS-SCHOLAR", in Proceedings of the National Computer Conference, San Diego, 1975, pp. 775-781.

Grignetti, M. & Warnock, E., Mixed-Initiative Information System for Computer-Aided Training and Decision Making, Electronic Systems Division TR-73-290, September 1973.

Kahn, K., An Actor-Based Computer Animation Language, MIT Artificial Intelligence Laboratory Working Paper No. 120 (LOGO Working Paper No. 48), February 1976.

Kahn, K., A Knowledge-Based Computer Animation System, MIT Artificial Intelligence Laboratory Working Paper No. 119 (LOGO Working Paper No. 47), February 1976.

Neches, R., On the Major Components of an Intelligent Educational Dialogue System, Muir College, November 1975.

Stansfield, J., Programming a Dialogue Teaching Situation, Unpublished PhD Thesis, Department of Artificial Intelligence, University of Edinburgh, Scotland, 1975.

Stansfield, J., B. Carr, & I. Goldstein, Wumpus Advisor I: A First Implementation of a Program that Tutors Logical and Probabilistic Reasoning Skills, MIT Artificial Intelligence Laboratory Memo No. 381, Sept. 1976.

Winston, P., Learning Structural Descriptions From Examples, MIT Artificial Intelligence Laboratory Technical Report No. 231, Sept. 1970.

Yob, G., "Hunt the Wumpus", Creative Computing, Sep-Oct, 1975, pp. 51-54.